

Proximity Displays for Access Control

Kami Vaniea

CMU-CS-12-141

September 2012

Computer Science Department
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Lujo Bauer, Co-chair

Lorrie Faith Cranor, Co-chair

Jeannette Wing

Michael K. Reiter, *University of North Carolina, Chapel Hill*

Stuart Schechter, *Microsoft Research*

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.*

© 2012 Kami Vaniea

This research was sponsored by the National Science Foundation under grant numbers CNS-0433540, CNS-0627513, CNS031428, CNS1116934, and DGE-0903659; and the U.S. Army Research Office under grants numbers DAAD190210389 and W911NF0910273.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

Report Documentation Page		Form Approved OMB No. 0704-0188
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.		
1. REPORT DATE SEP 2012	2. REPORT TYPE	3. DATES COVERED 00-00-2012 to 00-00-2012
4. TITLE AND SUBTITLE Proximity Displays for Access Control		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER	
	5e. TASK NUMBER	
	5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,School of Computer Science,Pittsburgh,PA,15213		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited		
13. SUPPLEMENTARY NOTES		

14. ABSTRACT

Managing access to shared digital information, such as photographs and documents, is difficult for end users who are accumulating an increasingly large and diverse collection of data that they want to share with others. Current policy-management solutions require a user to proactively seek out and open a separate policy-management interface when she wants to review or change her access-control policy. However, end users treat access control as a secondary task, and rarely visit a website for the primary task of managing security. Historically, security administrators and auditors were available to check for access-control issues on behalf of users, but in the age of Facebook and Flickr people are responsible for their own content. Users need a way to review their access-control policies that fits into their normal workflows. This thesis proposes the use of proximity information displays | small interface components spatially located near the data elements (or near a representation of data, e.g., the name in a file manager or thumbnail photo in a photo album) that contain information about who currently has access or who could access the data. These displays are intended to help users become more aware of how their data has been used in the past and how the data could be used in the future. We present empirical studies that test the hypothesis Users of a system that includes proximity information displays of access control-information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface. The focus of this thesis is understanding the impact of proximity displays on people's permission-modification behavior. The displays were conceptualized based on interviews with end users and security administrators, which highlighted the need for increased end-user awareness of their policies. Focus groups showed that people liked the idea of showing permission information in proximity to data. Finally, several evaluation studies were conducted in the lab and online using a photo-sharing website. Participants who saw proximity displays that were more comprehensive and could be glanced at easily were better able to identify access-control policy errors. Participants who saw displays that were overly coarse-grained, on the sidebar, or showed information about who had previously viewed the photos, showed no improvement over those who saw permission settings only on a secondary interface. Our studies suggest that proximity displays for access control can help significantly the majority of users who do not normally check their access-control policies.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:

a. REPORT
unclassified

b. ABSTRACT
unclassified

c. THIS PAGE
unclassified

17. LIMITATION OF ABSTRACT

**Same as
Report (SAR)**

18. NUMBER OF PAGES

278

19a. NAME OF RESPONSIBLE PERSON

Keywords: security; privacy; visualization; policy; usability; photo sharing; policy authoring; user interface; access control

This thesis is dedicated to

my family,
for their continuous encouragement,

and my friends,
who were a constant source of support and inspiration.

Abstract

Managing access to shared digital information, such as photographs and documents, is difficult for end users who are accumulating an increasingly large and diverse collection of data that they want to share with others. Current policy-management solutions require a user to proactively seek out and open a separate policy-management interface when she wants to review or change her access-control policy. However, end users treat access control as a secondary task, and rarely visit a website for the primary task of managing security. Historically, security administrators and auditors were available to check for access-control issues on behalf of users, but in the age of Facebook and Flickr people are responsible for their own content. Users need a way to review their access-control policies that fits into their normal workflows.

This thesis proposes the use of *proximity information displays* — small interface components spatially located near the data elements (or near a representation of data, e.g., file name in a file manager or thumbnail photo in a photo album) that contain information about who currently has access or who could access the data. These displays are intended to help users become more aware of how their data has been used in the past and how the data could be used in the future. We present empirical studies that test the hypothesis:

Users of a system that includes proximity information displays of access control-information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface.

The focus of this thesis is understanding the impact of proximity displays on people's permission-modification behavior. The displays were conceptualized based on interviews with end users and security administrators, which highlighted the need for increased end-user awareness of their policies. Focus groups showed that people liked the idea of showing permission information in proximity to data. Finally, several evaluation studies were conducted in the lab and online using a photo-sharing website. Participants who saw proximity displays that were more comprehensive and could be glanced at easily were better able to identify access-control policy errors. Participants who saw displays that were overly coarse-grained, on the sidebar, or showed information about who had previously viewed the photos, showed no improvement over those who saw permission settings only on a secondary interface. Our studies suggest that proximity displays for access control can help significantly the majority of users who do not normally check their access-control policies.

Acknowledgments

First and foremost I would like to thank my wonderful adviser Lorrie Faith Cranor for all the work she has invested in my education. Training new graduate students is not always easy, but she has shown nothing but patience and commitment.

As an undirected incoming student Michael K. Reiter and Lujo Bauer made me part of their research on the smartphone based distributed access-control system called Grey. Being part of a large project with students from a range of backgrounds and interests really helped me understand what conducting interdisciplinary research was like. It also helped me cement an interest in usable access-control technologies. I want to particularly thank Mike and Lujo for never accepting good enough work, and always demanding excellence. Frustrating as it was at the time, their lessons have made me a better researcher.

My thesis committee Lorrie Faith Cranor, Lujo Bauer, Michael K. Reiter, Stuart Schechter, and Jeannette Wing all deserve recognition for their tireless attention to detail when reading my thesis. Their comments provided much needed perspective. Writing a thesis is a joint experience and I owe all my thesis committee members a thank you for their assistance.

Being a member of the Cylab Usable Privacy and Security (CUPS) lab was an incredible part of working with Lorrie Faith Cranor. During my time at Carnegie Mellon I have had the pleasure of working with fourteen CUPS PhD students on a variety of projects. While every one of these students has helped me on my path to thesis completion, there are a few names that deserve to be mentioned. Rob Reeder's infectious personality and "that is an awesome idea" attitude really made even the most dull parts of research fun. Janice Tasi in her unofficial role as lab social coordinator made sure all of us showed up to parties occasionally. Rich Shay, Saranga Komanduri, and Manya Sleeper significantly assisted in my thesis work by answering endless grammar questions, and sounding out ideas.

The Computer Science Department staff members have been wonderful throughout my time as a student. Particularly Deborah Cavlovich and her predecessor Sharon Burks who helped me navigate the murky waters surrounding all the necessary forms required as part of being a graduate student. Not to be forgotten, are also Tiffany Todd, Jennifer Lucas, and Karen Lindenfelser who greatly helped me with any and all day-to-day administration challenges I brought to them.

My family has been nothing but supportive through all of graduate school. They have put up with many years of my physical absence, late birthday cards, and lack of returning phone calls.

I have made truly wonderful friends in grad school, as well as retained many from undergrad. All of them have provided vital emotional support, and contributed to my much needed social life. Every time I try and think of a small set of names to call out I realize how many incredible friends I have.

Finally, I would like to thank my two kittens Nala and Dash for posing for many of the photographs used in this thesis.

Contents

1	Introduction	1
1.1	Access-control proximity displays	2
1.2	Thesis statement	3
1.3	Research questions	3
1.4	Overview of studies	4
1.4.1	Reactions to access-control proximity display content	4
1.4.2	Proximity information display quantitative and qualitative evaluation	4
1.5	Outline of the thesis	6
2	Related work	7
2.1	Placing security information in spatial proximity	8
2.2	User policy management	8
2.2.1	User awareness	9
2.2.2	Policy reevaluation	9
2.2.3	Managing permissions in the home	9
2.2.4	Managing permissions in an organizational setting	10
2.2.5	The social statements access-control settings make	11
2.2.6	Automatically create policies	12
2.3	Enabling access control management	12
2.3.1	<i>Ex-ante</i> control	13
2.3.2	<i>In-medias-res</i> control	15
2.3.3	<i>Ex-post</i> control	16
2.4	Access-control policy tactics	17
2.5	Behavioral models	19
2.5.1	C-HIP model	19
2.5.2	HITL framework	20
3	Supporting end-user permission management	23
3.1	Potential research directions	24
3.1.1	Remove users from the loop	24
3.1.2	Pull: Wait for user's request	24
3.1.3	Push: Proactively show information	24
3.2	Chosen direction	25
3.3	Proximity displays	26

3.3.1	Scenarios	26
3.3.2	Design space	27
4	Focus group: User reactions to proximity security information	31
4.1	Interface designs	31
4.2	Methodology	33
4.2.1	Participants	34
4.2.2	Protocol	34
4.3	Results	34
4.3.1	Why is privacy important to me?	36
4.3.2	Who has viewed my photos?	37
4.3.3	Who could see my photos?	38
4.3.4	Proximity displays in personal and work environments	39
4.4	Design implications	39
4.5	Conclusion	40
5	Proximity access-control information displays	43
5.1	Design space	43
5.2	Platform	44
5.3	Proximity display plug-in design	45
5.3.1	Implemented policy shown in a grid-based design	45
5.3.2	Implemented policy shown in a list-based design	46
5.3.3	Audit information shown in a list-based design	47
5.3.4	Facebook icons	49
5.4	Access-control permission-modification interface	50
5.4.1	Full-page interface	50
5.4.2	Dialog interface	50
5.4.3	Conflict resolution and effective permissions	53
6	Detailed methodologies	55
6.1	Hypotheses	56
6.2	Eye-tracker study	56
6.2.1	Study conditions	56
6.2.2	Protocol	57
6.2.3	Recruitment and demographics	59
6.2.4	Data collection and analysis	59
6.3	Lab study	60
6.3.1	Study conditions	60
6.3.2	Protocol	63
6.3.3	Participants	66
6.3.4	Data collection and analysis	67
6.4	Online study	70
6.4.1	Study conditions	70
6.4.2	Participants	72

6.4.3	Protocol	72
6.4.4	Data analysis	78
6.5	Conclusion	81
7	Effectiveness of proximity displays	83
7.1	Hypothesis testing	84
7.1.1	H1: Correcting/checking permissions	85
7.1.2	H2: Permission recall	85
7.1.3	H3: Negative effects	87
7.2	How people notice and fix permission errors	87
7.2.1	Noticing permissions	88
7.2.2	Participants' tendency to check permissions	92
7.2.3	When do people change permissions	95
7.3	Proximity-display designs	103
7.3.1	Under photo	104
7.3.2	Sidebar	104
7.3.3	Mixed	104
7.3.4	Facebook	105
7.3.5	Audit	105
7.4	Limitations	106
7.5	Conclusion	107
8	Designing an access-control study where security is a secondary task	109
8.1	Study goals	110
8.2	General study design	111
8.3	Secondary permission task	114
8.3.1	Study 1	114
8.3.2	Study 2	116
8.3.3	Study 3	116
8.3.4	Study 4	117
8.4	Participant responsibility	118
8.4.1	Study 1	118
8.4.2	Study 2	119
8.4.3	Study 3	119
8.4.4	Study 4	120
8.5	Ideal policy comprehension	120
8.5.1	Study 1	120
8.5.2	Study 2	121
8.5.3	Study 3	122
8.5.4	Study 4	122
8.6	Effective outcome measurement	123
8.6.1	Study 1	123
8.6.2	Study 2	123
8.6.3	Study 3	124

8.6.4	Study 4	126
8.7	Discussion	127
9	Conclusion	129
9.1	Contributions	130
9.2	Future work	131
9.2.1	Proximity-display design	131
9.2.2	Understanding policy error-identification behavior	132
9.2.3	Exploring proximity displays in other domains	132
10	Bibliography	135
A	Focus group study	145
A.1	Focus Group Script	145
A.1.1	Information visualization explanations	148
A.2	Focus group 1 packet	150
A.3	Focus group 2 packet	159
A.4	Focus group 3 packet	167
A.5	Focus group 4 and 5 packet	176
B	Eye-tracker study (study 2)	185
B.1	Printed instructions and emails	185
B.2	Online survey	205
C	Lab study (study 3)	217
C.1	Printed instructions and emails	217
D	Online study (study 4)	235
D.1	Online survey	235
E	Evaluation data	245

List of Figures

1.1	Proximity display showing access-control settings under an album thumbnail.	2
2.1	Communication-Human Information Processing Model (C-HIP).	19
2.2	Human In The Loop Framework	21
4.1	Example interface typical of the ones shown to focus group participants. This interface shows detailed implemented-policy information and summarized audit information adjacent to the album thumbnails.	33
4.2	Usage scenarios illustrating how an end user might use proximity displays both to cause a positive social experience and to notice an issue.	35
5.1	The proximity displays shown to users in the four evaluation studies. The display used in the first and second studies (a), is based on a grid-style design. The displays used in the third (b) and fourth (c) use a list-based design. Displays (a) and (b) include a “Manage Permissions” link; participants were rarely observed to use the link, so it was removed in design (c).	46
5.2	The proximity displays showing who had accessed the album (audit). Unlike the displays shown in Figure 5.1, which show who could access the album in the future, these displays show who has accessed the album in the past. Figure (a) was pilot tested during evaluation studies 2 and 3, resulting in the Figure (b) design, which was evaluated in the final evaluation study (study 4).	48
5.3	Full-page policy-modification interface used by participants to make changes to the access-control policy. All the albums are listed along the left; user groups are listed along the top of the grid; and view, edit, and add permissions are shown as icons in the central grid. This interface also contains a legend at the bottom left.	51
5.4	Permission-modification dialog. Sentence at the top of the dialog reminds users what album the permissions refer to. The group names are listed along the left side, followed by the different actions (view, edit, and add) that are allowed or denied for that group. A black icon indicates that the permission is allowed; a light grey icon indicates that the permission is denied. Placing the mouse over any icon produces a tool tip indicating the meaning of the current icon. For example: “Animal Shelter can view this album.” Clicking on an icon toggles it between allow and deny.	52

6.1	The two proximity display conditions used in the eye-tracker study: (a) in the under-photo condition, and (b) the sidebar condition. Proximity display in the under-photo condition. The proximity display in (a) shows that the group Everybody has no permission; Coworkers can view and add to this album and all subalbums, but can only edit some subalbums; Family can view this album and some subalbums; and Friends cannot view anything. .	57
6.2	Gallery 3 interface without a proximity display (a), and with a proximity display under every photo and album (b).	61
6.3	Full-page permission-modification interface (a) and dialog permission-modification interface (b).	62
6.4	Lab study protocol order.	63
6.5	Gerald's Photograph Policy	64
6.6	Participants were asked, by a co-worker, whether each of the above images were acceptable to post on Gallery 3 and whether the co-worker needed to make sure to do anything when they put the photo on Gallery 3. Q1 has no problems with the photo, but should be visible to Friends and Coworkers only. Q2 needs to be rotated and should be visible to Everybody on the Internet. Q3 cannot be uploaded to Gallery 3 because it is blurry and contains alcohol.	65
6.7	Gallery 3 interface showing all the albums and their cover thumbnails (a), and the interface showing all the photos contained within a single album (b). Proximity-display locations are marked with numbers 1-6 indicating the different locations where proximity displays were tested.	74
6.8	Screenshot from the online study showing the instructions and website frames. A control bar at the bottom of the instruction frame allowed participants to shrink the frame, obtain instruction on Gallery 3's features, and move to the next task.	77
6.9	Ideal policy rules in the online study.	77
6.10	Sample permission recall question from the post-survey. The question asks participants to recall both what the permissions should have been and what the permissions were at the end of the study.	80
7.1	Histogram of the number of fixations on proximity displays for all participants (y-axis) against the amount of time spent on the page, normalized (x-axis). Data is from the eye-tracker study.	89
7.2	Histogram of the difference between the number of tasks on which participants checked permissions when the tasks had an error and the number of tasks on which they checked when the tasks did not have an error. For example, in the lab study 11 control-condition participants checked the same number of permissions in tasks with errors as they did in tasks without errors. In the lab study (Subfigure (a)) "checked" was defined based on behavior observed by the researcher administering the study. In the online study (Subfigure (b)) "checked" was defined as opening the permission-modification interface.	90

7.3	Number of tasks where the permission-modification interface was opened by participants in the control and under-photo conditions. Graph <i>a</i> shows data from the lab study, and graph <i>b</i> shows data from the online study. Graphs of other conditions in the online study are nearly identical.	93
7.4	While working on tasks in the lab study, participants were free to engage in actions in any order, including interleaving actions. For example: a participant could rotate a photo, delete a photo, then rotate a photo. Graph <i>a</i> shows the first time an action of that type was engaged in during a particular task and whether that action was the first action, the last, neither first nor last (middle), or the only action engaged in. The height of the bars indicates the total number of tasks across all users; the summation of all bars in a subgraph is the number of tasks, across all users, in which the action was engaged in at least once. Graph <i>b</i> is similar to graph <i>a</i> except that it shows the last time a action is engaged in during a task.	97
7.5	The number of seconds into the task when the permission-modification interface was opened by participants in each condition. Events from task 1 and the training are excluded to remove bias caused by prompting participants.	99
7.6	As part of the lab-study verbal post-survey, participants were asked to recall Gerald's rules, in their own words. The above graph shows the order in which participants recalled the rules. The majority of participants recalled the rules in the following order: R1, R5, R2, R3 and forgot to mention R4.	101
7.7	Number of seconds into a task that an action was engaged in (10 second intervals). Histograms show all participants across all conditions, both with and without permission proximity displays. Events from task 1 and the training are excluded to remove bias caused by prompting participants. .	102
8.1	Example of proximity display used in studies 1 and 2. The interface for studies 3 and 4 had a slightly different permission display interface design.	112
8.2	Email from Pat's friend implying that the Friends group needs to be able to view the photos.	115

List of Tables

- 6.1 Tasks and information given to eye-tracker study participants. 58
- 6.2 Implemented policy errors. Each participant experienced every error once during the 14 tasks. The first training task had a permissions error where everybody on the Internet could see a personal album, so participants would have seen this error twice during the study session. 68
- 6.3 Position and type of access-control proximity display shown for each condition and page. 75
- 6.4 Position and type of tag proximity display shown for each condition and page. 76
- 6.5 The number of users in the online study who completed the study in each condition, the number of participants who made at least one change to permissions in either condition, and the number of participants who agreed to the consent form but did not complete the study. 78
- 6.6 Tasks and their associated permission and tag errors. 78
- 7.1 Methodologies used in each study. 84
- 7.2 The conditions tested in each study; details on each condition can be found in Section 6.4.1. 84
- 7.3 Results of Wilcoxon signed-rank statistical test (within-subjects). We present both the median and the average number of permissions corrected in the control and experimental conditions. 85
- 7.4 The online-study participants’ ability to recall permission settings for two non-training albums (five questions each). Reported p-values reflect Holm-Bonferroni correction. 86
- 7.5 The online-study participants’ ability to apply the permission rules in the ideal policy for the two non-training albums per condition (5 questions each). Participants were asked what permissions Pat/Pat’s boss would have wanted to set. Reported p-values reflect Holm-Bonferroni correction. 87
- 7.6 In addition to tag and permission errors, the online-study participants were asked to correct issues with the titles, organization, orientation, and content of photos. This table reports the number of non-permission and non-tag errors participants corrected out of 37 errors. Reported p-values reflect Holm-Bonferroni correction. 87
- 8.1 Order, name, and methodology of each study. 109

E.1	Data depicted in Figure 7.1(a).	245
E.2	Data depicted in Figure 7.1(b).	246
E.3	Data depicted in Figure 7.2(a).	246
E.4	Data depicted in Figure 7.2(b).	247
E.5	Data depicted in Figure 7.3(a).	247
E.6	Data depicted in Figure 7.3(b).	248
E.7	Data depicted in Figure 7.4 a.	249
E.8	Data depicted in Figure 7.4 b.	250
E.9	Data depicted in Figure 7.5 (sidebar).	251
E.10	Data depicted in Figure 7.5 (under).	252
E.11	Data depicted in Figure 7.5 (mixed).	253
E.12	Data depicted in Figure 7.5 (facebook).	254
E.13	Data depicted in Figure 7.5 (audit).	255
E.14	Data depicted in Figure 7.6.	256
E.15	Data depicted in Figure 7.7.	257
E.16	Data depicted in Figure 7.7 cont.	258
E.17	Data depicted in Figure 7.7 cont.	259

Chapter 1

Introduction

End users find it challenging to stay aware of and manage sharing preferences for content that they publish on social networks and photo-sharing sites [16, 65, 104, 106]. This problem is becoming even more difficult as sites become more dynamic, with constant uploading of content and shifting groups of friends. In this dynamic environment, security policies are difficult not only to set up, but also to maintain. When the current implemented policy changes due to a new group member or the user's social interactions with a group member, mismatches in the implemented policy can occur.

Having a mismatch between the currently implemented access-control policy and the policy users believe to be enacted can place end users in dangerous or awkward situations. If we turn to the news, we see numerous accounts of users who set their permissions incorrectly and experienced a loss because of it. A girl in Germany accidentally publicized her birthday party on Facebook and ended up with 15,000 RSVPs, 1,600 of whom actually showed up [61]. A teaching student was denied her diploma after a photo of her drinking was shared publicly online [58].

In an ideal world the computer system would analyze user behavior and continuously maintain the access-control policy, dealing with changing environments and preferences. Unfortunately, computer policy management systems can never be perfect. While systems do exist that will detect and flag potential issues with access-control policies, those systems are limited by their understanding of what users currently want their policy to look like, a preference that can change frequently. Without this knowledge, policy error detection systems are prone both to missing important errors and to flagging policy components that are error free.

Because programmatically creating and maintaining access-control policy with high accuracy is not currently feasible, it falls to the end user to periodically check and adjust their policy to meet their current needs. End users' sharing intentions change over time as their social environment evolves, and the content being protected changes. Additionally, sites such as Facebook periodically add and remove privacy/access-control settings, effectively altering the access-control policy on behalf of users. The end result is that even if users correctly implement their intended policy using available settings, that policy would likely develop errors over time.

Internet users claim that security and privacy are important to them but the reality is that Internet users rarely interact with access-control policy as their primary task [25]. They log onto Facebook, Flickr, or Google+ to share content, catch up on news, and interact with friends—not to “do security.” Access-control typically remains in the background until some event, such as an embarrassing experience, brings it to the user’s attention [30, 104]. So users are unlikely to identify errors in their current settings unless they actively decide to look for them.

Providing end users with usable privacy controls is starting to be seen as a marketable feature by websites built on user content. Social networking sites such as Google+ are trying to differentiate themselves from their competitors by providing users with more usable privacy controls. In recent years we have seen these sites move away from placing all the privacy settings on a secondary page, and start putting some of them near the data element they control. However, to my knowledge, the effectiveness of these displays has not been empirically tested.

In prior work, I and others have studied how people interact with access-control technology [9, 10, 11, 29, 57, 69, 79, 94]. That research has yielded a better understanding of the issues people and organizations have with managing access control. However, one of the most striking issues for me was the observation that many users have mismatches between the access-control decisions computer systems are making on the user’s behalf and what decisions users would like the computer systems to make. I considered several different ways to help users identify and correct these mismatches. I ultimately decided to try using proximity information displays to convey the current permission settings and to convey access events to end users.

1.1 Access-control proximity displays

In this thesis I am proposing the use of proximity information displays to make users more aware of how their resources have been used in the past and how they could be used in the future. Proximity information displays (Figure 1.1) are interface components that show users information about their access-control settings and who has been accessing their resources in a way that is easy to understand and enable users to create policies that better

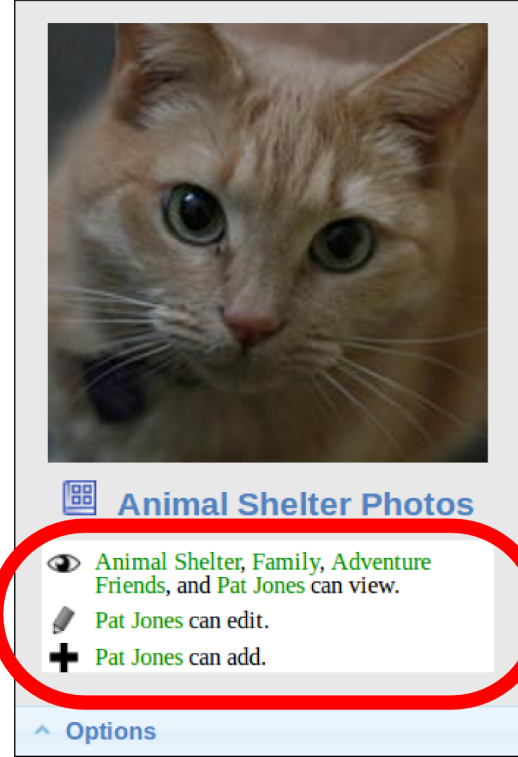


Figure 1.1: Proximity display showing access-control settings under an album thumbnail.

match their preferences. They are referred to as *proximity* information displays because information is always placed in close proximity to places where users interact with or think about their resources.

Proximity displays are designed to enable users to notice permissions. They should enable users 1) to better understand the current access-control settings, and 2) to identify mismatches between what they want and the current settings.

Egelman divides the space of security indicators into *passive* and *active* [32]. Active indicators force users to make a decision before progressing. Passive indicators present information to users but do not force users to notice or engage with the indicator. Proximity displays are intended to be passive indicators. While many end users' access-control policies do not necessarily match their access-control preferences, programmatically identifying access-control policy mismatches is error prone with potentially high false-positive rates. Proximity displays are intentionally designed to be passive. The end user should be able to notice permission errors easily, while not experiencing a negative impact to their normal workflow.

In this thesis work I explore the design of proximity information displays and the effect the displays have on users' ability to identify issues with, and be aware of their access-control policies.

1.2 Thesis statement

The objective of this thesis is to test the hypothesis:

Users of a system that includes proximity information displays of access-control information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface.

1.3 Research questions

In this thesis I take an in-depth look at how people notice access-control errors and the impact proximity access-control displays have on that behavior. My work addresses a range of questions intended to support my thesis topic. These questions, enumerated below, express the specific issues I will be looking at in this thesis.

1. How do people react to access-control setting information being presented on the same screen as their photos? Chapter 4
2. How do people react to information about who has previously interacted with their photos being presented on the same screen as their photos? Chapter 4
3. What is an effective lab environment design that enables participants to both understand the goal and still treat security as a secondary task? Chapter 8
4. Do proximity displays improve people's ability to identify access-control permission errors over having the information on a secondary screen? Chapter 7

5. Do proximity displays improve people’s ability to remember their access-control permission settings? Chapter 7
6. Do proximity displays negatively impact people’s performance on their primary task compared to having the information on a secondary screen? Chapter 7
7. Do proximity displays which show who has previously interacted with the photos improve peoples’ ability to identify access-control permission errors over having setting information on a secondary screen? Chapter 7
8. Does the position of the proximity display impact people’s ability to notice errors? Chapter 7

1.4 Overview of studies

I present the results of five studies we conducted in order to examine how people react to access-control information placed in close spatial proximity to the item it controls.

1.4.1 Reactions to access-control proximity display content

To better understand how people would react to different types of information and different presentation methods I conducted a focus group study. The interviews suggested that people had need of a detailed display that gave them concrete information on which to base their mental models of their security policy. The interviews also suggested that presenting detailed information about who had previously accessed their photos would assist users in their continued reevaluation of their policies and social networks. However, participants considered detailed information about who had viewed their photos to be highly invasive because it “forced [them] to stalk [their] friends.” Participants were generally positive about showing setting information, provided that it did not take up too much screen real estate. Several users commented about the positive effect of finding and changing permissions easily. In Chapter 4 I discuss the high-level take aways from the focus groups.

1.4.2 Proximity information display quantitative and qualitative evaluation

The positive view of focus group participants suggests that proximity information displays that show permission setting information are perceived as useful. However, I wanted to know if these displays are actually useful for participants in terms of assisting them 1) to identify errors in their policies, and 2) to improve their awareness of the content of their policies. To test the actual usefulness of the displays I conducted several role-play lab studies where I asked participants to come into the lab and work through several tasks while playing the role of a fictitious person who managed an online photo-sharing site. This person was responsible for fixing permission and non-permission errors, such as spelling, orientation, and tags. Participants were told what the access-control policy should be for different types of photos in the albums. They were then given a series of emails that

requested various changes to the photo albums. In the course of fulfilling these requests they had the opportunity to detect and correct permission errors. I conducted four lab studies using this format to quantitatively and qualitatively understand the effect proximity displays have on participants' permission error identification and policy awareness behavior. In Chapter 6 I detail the methodologies used to test the proximity displays, and in Chapter 7 I detail the combined results of these studies.

Study 1, pre-study: Based on the responses to proximity displays in the prior studies, we decided to focus on presenting permission information on the proximity displays and leave the information about who has seen the photos for a later study. Based on the focus group feedback, I was concerned about the amount of screen real estate required by the proximity display. I was also concerned about the effect of putting the display in an obscure a location. To evaluate these concerns, I tested the display both on the sidebar and under every photo and album thumbnail. The outcome of this study was inconclusive and the study was stopped early due to several methodological issues (Chapter 8), but the behaviors of the participants strongly indicated that showing permission information in close spatial proximity enabled participants to notice errors in their permission settings.

Study 2, eye-tracker study: In the pre-study I observed a participant behavior I term "checklisting." Participants who checklisted would appear to finish with a task, pause, go through a check list of all the types of actions we had trained them on, and then explicitly check the permissions. The methodology from the pre-study was redesigned to reduce this behavior by reducing the number of error types, both permission and non-permission, present in each task. I also added some qualitative data collection mechanisms, including an eye tracker, to better capture how participants were interacting with the proximity displays. The result of this study was that placing proximity displays under every album thumbnail and photo enabled participants to identify statistically significantly more errors than placing it on the sidebar or placing access-control setting information on a secondary page. I also observed that participants who see proximity displays under the photos tend to see the displays mid-way through the task, but change the permissions at the end of the task.

Study 3, lab study: In the prior study I saw a statistically significant difference in the number of permission errors identified, but I did not see a difference in participants' ability to remember the permissions. Results from the interview study indicated that participants reason about their security settings off line, and make decisions that depend on their memory of their settings being correct. In addition to enabling participants to find permission errors, I also wanted to make them more aware of their current settings. I hypothesized that the lack of difference in memory in the prior study was caused by 1) forcing control participants to repeatedly access the permission modification interface and, 2) providing participants with a permission modification interface that showed the policies for all albums, not just the album participants are currently working with. In this study I decided to test the style of the permission modification interface used, in addition to the proximity display. I also increased the amount of qualitative data collection by adding a post-study interview where I used a cognitive interview approach to ask participants about the choices they made during the study. I found that the permission-modification interface used impacts participants' ability to notice errors, but had no impact on memory. I also

learned that participants were able to glance at the displays and some participants had a natural tendency to correct all permission errors in one single pass.

Study 4, online study: The prior studies indicate that proximity displays help people identify permission errors. However, these studies were done with a small number of participants. The results of the prior study also highlighted the high level of participant variability; some participants are more inclined to check permissions than other participants. To address this, I conducted a within-subjects study — every participant saw the control condition and one of the proximity-display conditions. I also increased the number of proximity-display conditions including proximity-display designs that mimic the Facebook proximity-display design and a proximity display that contains information about who has seen the photo album. I found that conditions that used proximity displays that showed permission setting information under photos/albums or under album thumbnails and on the sidebar were statistically significantly better than control at enabling participants to identify errors. However, similar to the prior studies, I saw no difference in memory.

1.5 Outline of the thesis

In this thesis I start with a discussion of the related work (Chapter 2), discussing both what is currently known about the way people manage access control, and systems similar to proximity displays. Then I motivate the need for proximity displays (Chapter 3). This is followed by a focus group study to explore peoples’ reactions to variations in proximity display content and design. The details of the proximity display design and implementation in the Gallery 3 photo sharing system are described in Chapter 5. I detail the methodologies of the last three studies (Chapter 6), then describe the results (Chapter 7). Designing the methodology for the four studies that tested the effectiveness of the proximity display was an informative experience with several lessons learned (Chapter 8). Finally, I conclude with a discussion of the contributions, future work, and design recommendations for proximity displays (Chapter 9).

Chapter 2

Related work

Making the process of managing access control usable is a difficult and important problem. In 2003 the Computing Research Association released a list of four Grand Research Challenges including: “Give end-users security controls they can understand and privacy they can control for the dynamic, pervasive computing environments of the future” [46].

“Security features in IT systems are, in a sense, like brakes on automobiles. Although brakes are used to slow or stop vehicles, their real purpose is to enable drivers to go faster by enabling them to avoid accidents caused by external threats (such as mechanical failure in other vehicles, rude or reckless drivers, road hazards, stop signals and heavy traffic). Better security is an enabler for greater freedom and confidence in the cyber world” [46].

The National Academy of Engineering agrees, listing cybersecurity as one of their Grand Challenges and specifically noting that understanding the psychology of computer users is a vital component of improving the state of cybersecurity in general [72].

Systems that allow end users to configure privacy settings may be thought of as access-control or security systems, as they involve policies that govern access to a user or to a user’s personal information. In this thesis we will be using the terms *privacy settings* and *access-control settings* interchangeably to refer to the set of settings users can manipulate to control who can see what part of their information. The term *implemented* is used to refer to the current state of the access-control settings on a system. The term *ideal* refers to the access-control policy users would like to implement on the system.

In this chapter we will discuss what the research community currently knows about how end users manage implemented and ideal access-control policies. End users currently have difficulty managing their implemented policies, which results in negative consequences for users (Section 2.2). Researchers are designing systems that better support users in their access-control management tasks (Section 2.3). Finally, there exist models that explain end-user behavior towards warnings in general and warnings specific to computer security (Section 2.5).

2.1 Placing security information in spatial proximity

Several studies have looked at the effect of showing privacy- and security-related information in proximity to items users are working with [32, 62, 98, 99, 103]. These studies generally show that displaying security and privacy information in proximity to related items can positively influence end-user behavior provided users understand the import of the information.

Tsai et al. showed that placing a graphical representation of each website’s privacy policy next to search results increased the amount of money people were willing to pay when purchasing privacy sensitive items [98, 99].

Lieberman et al. were concerned about the impact of accidentally emailing too many people through lists or the “reply to all” feature [62]. They designed an email interface to enable people to easily spot errors. The interface showed the photos of all the people being emailed near the box for email addresses. They saw a significant improvement in people’s ability to quickly identify who was being emailed.

Wang, in his thesis work, displayed a large privacy related proximity display on the side of a fictional book selling website [103]. He observed that participants self reported interacting with the display and liking the options.

Egelman divides the space of security indicators into *passive* and *active* [32]. Active indicators force users to make a decision before progressing. Passive indicators present information to users but do not force users to notice or engage with the indicator. Sunshine et al. showed that passive indicators were less effective than active indicators in alerting users to the dangers of self signed certificates [96]. Sotirakopoulos et al. repeated the Sunshine study and found that users ignored both active and passive indicators [91]. The authors hypothesize that users were now more familiar with the active indicators and were now habituated to ignoring them.

Kelly et al. attempted to use eye tracking to better understand how people, particularly experts versus non-experts, looked at Facebook’s proximity display icons while normally interacting with their own Facebook profiles [55]. Their study was plagued by issues related to the eye tracker technology used. However, their preliminary results show that some users do look at the access-control information.

2.2 User policy management

End users have trouble managing permissions in their online social environments. They are not aware of what their current permission settings are, and they incorrectly assume permissions to be correct when errors exist. They do not regularly check the permissions for errors or adjust their policies when their context changes. This can lead to a loss of privacy for individuals, resulting in potential harm and embarrassment. The goal of this thesis is to help people maintain access control policies that match their preferences. In this section we take a look at how people currently interact with security at home and at work.

2.2.1 User awareness

Users often lack awareness of their current policies while inaccurately assuming that the policies are fine. An empirical study of Facebook users compared participants' sharing intentions to the implemented privacy policy, and found that every participant they interviewed had at least one mismatch [65]. A survey of Facebook users' understanding of applications found that only one out of 516 surveyed users was able to accurately answer what parts of their Facebook profile the survey application could access [56].

Whalen et al. conducted an online survey on end-user experiences with sharing and access control. They found that users have dynamic access-control needs that vary with different tasks and are often frustrated by current access-control mechanisms that are difficult to use and not well-suited to users' workflow [106].

Research in the area of photo access-control management shows that end users care about the privacy of their photos. Similar to other data-sharing domains, end users claim to care about privacy but have difficulty managing it. Ahern et al. [5] found that sharing decisions are often related to the people in photos and the environment in which they are taken. Besmer and Lipford [16] also report that "impression management" is a common concern when sharing photos.

2.2.2 Policy reevaluation

As part of interacting with others, people continuously construct, interpret and reevaluate their social context based on actions others take [73, 77]. In file security the actions of others are often hidden by the system, and even the settings are placed on secondary pages where they are not readily visible. Without this visibility it can be challenging for users to take the access-control permission settings into account as part of their natural social reevaluation [14]. Users may not realize that their access-control policy no longer accurately represents what they want until something happens to bring it to their attention [30]. Prior work in domains such as location tracking and photo management tells us that end users' policies can be very dynamic and are often based on the current relationships between users and the requesters [48, 57, 69, 74, 92]. However, when users set static policies to dictate who can see their location, unanticipated or out-of-character requests for location can be denied [23]. Studies of home and cloud file storage environments also show that end users start out by creating one access-control policy and, based on observations of how it is used, they may realize that an alternative policy is more appropriate [64, 68, 85].

2.2.3 Managing permissions in the home

Home users tend to view their immediate surroundings as non-malicious [68]. When working in small groups, people establish social rules that allow them to function without tight security. These rules work as long as the group is small but break down in larger settings [7]. Home users often trust the other members of their home and expect them not to pry beyond clearly marked boundaries. Instead of using technology to protect their files, users hide the files or put them on clearly marked personal devices or in personal

spaces [68, 84]. People tend to think about physical and virtual security holistically, not separating the two concepts [30, 68].

The home is rarely a structured environment, where each person has her own account and all her data and accesses are tightly tied to the account. Home users tend to share a single account on the “family computer,” preventing a clear tie between account and person [21, 84]. Account sharing is primarily driven by convenience issues, being able to quickly access the computer outweighs the privacy and security concerns that having multiple accounts would solve [33]. Users also have dynamic access-control policies that can change quickly [10, 63]. Using focus groups on ubiquitous shopping technology, Little, Silence and Briggs found that users are concerned that computer-controlled devices cannot properly respond to the unpredictable day-to-day behavior changes of the home environment [63].

Users are not necessarily skilled at managing their computing resources on their own and tend to seek help from trusted people when they need it [30, 75]. Once a trusted person has been consulted, users tend to blindly believe that the device is now secure even if the trusted person is no longer present [30]. Users also appear to learn about “correct” security behavior from stories told to them by other users. These stories allow users to learn from the negative experiences of others [77]. Users create heuristic rules about the types of data that are stored on different devices and therefore accessible to different sets of people. These rules are rarely if ever updated [85].

2.2.4 Managing permissions in an organizational setting

Organizations maintain internal file systems that track important documents, preserve confidentiality, and ensure security protections but, traditionally, do not encourage end users to engage in secure sharing behavior. Many of these systems prevent internal organization employees from reliably sharing files with others and themselves [106]. The lack of reliable file-sharing support forces users to circumvent the perceived pointless impediment of the file system and turn to alternative file-sharing technologies such as email, instant messaging (IM), third-party storage (Dropbox), and USB drives [28, 70]. These alternative sharing mechanisms enable users to quickly and easily share the document with whomever they want, but they lack many of the security properties of the original file system.

Unlike the home environment, where the number of users is small and the assumption of non-malicious users may be reasonable, the office environment can be large and contain malicious users. Schneier writes [87]: “Access control is difficult in an organizational setting. On one hand, every employee needs enough access to do his job. On the other hand, every time you give an employee more access, there’s more risk: he could abuse that access, or lose information he has access to, or be socially engineered into giving that access to a malfeasant. So a smart, risk-conscious organization will give each employee the exact level of access he needs to do his job, and no more.”

Malicious employees are a major concern for organizations. A study by CERT of 49 insider attacks found that 59% of the “insiders” were former employees and 43% still had authorized system access at the time of the attack [6]. A TELUS Security Labs study found that 33% of security breaches reported in 2009 were due to insiders. Insider breaches were

reported by 17% of Canadian organizations.

Not all insider attacks involve data breaches. Some insider attacks are just people making use of resources provided by the organization in ways that violate the organization's rules. Dwayne F. Cross, a government worker, was convicted of computer crimes for looking at over 150 passport files. His reason: curiosity [41].

Another challenge is the mismatch in concerns and goals of security professionals and end users. Members of the information technology (IT) field often perceive end users as insecure and the cause of many security incidents [51]. End users often view security professionals and even their own IT departments as being overly paranoid and generally getting in the way of the work end users need to get done [4, 42, 45], and end users are sometimes correct in assuming that dealing with extra security tasks is not worth their time [45].

A few studies have surveyed needs for access-control systems from a holistic organizational perspective. Ferraiolo et al. studied the access-control needs of 28 commercial and government organizations and identified seven access-control approaches. One approach they discuss is *discretionary access control* (DAC), in which access is assigned to individuals and groups, who in turn may delegate that access to others. The authors note that DAC is well suited for organizations where end users have rapidly changing information access needs and must be able to specify access-control policy for resources they control. They also mention that the DAC approach is not suitable for organizations concerned with maintaining tight controls on access rights [37]. The introduction of Role-Based Access Control (RBAC) [36] was partially intended to address this issue by making the setting of access-control permissions better fit how organizations actually manage their settings.

2.2.5 The social statements access-control settings make

Smetters and Good examined the acquired access-control rights of employees in a large office environment. They found that access rights tended to be collected over time at the company and treated as a status symbol [90]. Sinclair et al. observed a large financial institution during an entitlement review of its employees' current permissions to resources within the company including file permissions. As part of the review, auditors asked employees to review their own access to files and applications and remove permissions to resources they did not actually need. Employees voluntarily removed 15% of their own access permissions because they "just didn't want to worry about having access to applications they didn't need" [89].

In addition to access-control being viewed as a status symbol, too much focus on keeping things secure can be viewed as paranoia. Gaw et al. studied a non-profit organization where maintaining security was an important part of employees' job descriptions [38]. They found that employees used secure communications only for important documents and not for other types of communications. This was partially because doing things like encrypting all email was perceived as paranoid.

2.2.6 Automatically create policies

Theoretically, the best way to assist end users in their permission modification is to automate the problem away. If computers could automatically determine the correct policy and just enact it with a high degree of accuracy, our problems would be over. Unfortunately, we do not yet have a system capable of reliably predicting the correct access-control permissions and enacting them on our behalf. Researchers have endeavored to study and predict our access-control preferences [27, 57, 83].

Fischbein et al. found that users' preferences for sharing or hiding their location information varied across time even when the location, time of day, and requester remained the same [15]. They hypothesized that the changes were due to contextual factors not visible to the system. Sadeh et al. found that end users were able to specify policies that matched their *ex-post* preferences only 79% of the time [83]. Cranshaw et al. used machine learning with user feedback and was only able to accurately match end users' *ex-post* preferences at best 87% of the time [27].

There have also been attempts to use an Attribute-Based Access Control model [49] to automatically create rules based on pre-existing attributes. Klemperer et al. explored the use of photo tags, combined with user-specified rules, to manage access-control policy for photos [57]. They found that using organizational type tags resulted in 27% of photos being erroneously marked as allowed or denied for at least one "friend," though only 7.8% of friend and photo combinations were erroneously allowed or denied access.

Researchers generally view full automation of access-control policy creation as unlikely to happen soon. This is partially due to the high false-positive rates described above, but also because of the issue of exceptions and emergencies. As Rissianen et al. says, there can be many different situations in which an access request could be made and only some of those situations are possible to anticipate [82]. As Edwards et al. points out, automating security enforcement may not always be beneficial [31]. Other researchers have similarly observed that users plan ahead for exceptional or unanticipated situations and need an access-control system capable of supporting this type of forward thinking and planning [9, 11, 20, 28, 76, 93, 95].

If we accept that full automation is unlikely to happen in the near future, then we must rely on end users to be actively involved in the creation and maintenance of their own access-control settings. Users need computer systems that enable them to manage security as part of their workflow. Researchers have proposed several systems that enable users to manage security as part of their normal system interaction.

2.3 Enabling access control management

In access-control literature we tend to think about access-control policy specification as a user task and policy enforcement as a computer task. Stevens and Wulf coin the term *Computer-Supported Access Control*, or CSAC, to emphasize that the technological mechanisms behind access control are only one part of how access control as a whole is practiced. They argue that access control should be designed as a supporting system where humans

work with computers to actively manage how files are accessed, not an automation system, where the computer automatically enforces policy without additional input from users [94].

In his work, Lampson described the task of setting access-control permissions in terms of proactive permission setting and automated enforcement. In other words, he assumed that permissions would be set before any access attempts and that the system would be solely responsible for judging the appropriateness of the request and enforcing it based on previously expressed access-control lists [60]. These assumptions that access control is set before the access and that enforcement should be automated are common in the security community [36, 39, 47, 67, 86].

Stevens and Wulf [94] and other researchers [9, 76, 82] have postulated that access-control management tasks are actually conducted in one of three ways. *Ex-ante control* is when the resource owner sets the policy before any anticipated accesses occur and the computer enforces it at the time of the access. *In-medias-res control* is where the access permissions are defined by the resource owner at the time of the access attempt. Finally, *ex-post control* is where the computer automatically grants access and the legitimacy of an access request is evaluated by the resource owner after the access has already taken place.

Researchers have looked at many different ways to assist users with their permission-modification tasks. While there are many different ways to assist users, the approach taken by researchers depends largely on how they assume users will interact with their technology. In this section we discuss different solutions proposed by researchers to address users who manage their access-control policy *ex-ante*, *in-medias-res*, and *ex-post*. Throughout the thesis we make use of these concepts while designing interfaces and while interpreting results.

2.3.1 *Ex-ante control*

Users engaging in *ex-ante* control create their access-control policies proactively, in advance of any access attempt, and based on how they anticipate the resource will be used in the future. The policies are then automatically enforced by the computer system that is responsible for interpreting the policy expressed by users based on the current context.

Traditionally, end users interested in engaging in *ex-ante* control proactively seek out an access-control management interface and use it to specify the policy. In Windows XP, for example, users must proactively right click on a file and select the “Sharing and Security” option before they can see or modify the file’s policy. Many different researchers have looked at how to support this type of access-control policy management [19, 78, 80, 101].

Johnson et al. [50] built a system where end users could, through their email client, upload a document to a document sharing system, automatically grant access to all email recipients, and include a link to the document in the email instead of the document itself. Though they were never able to get the system to a fully deployed state, the researchers were able to observe participants’ positive reactions to it. As a result of this work they put forward the idea of Laissez-faire access-control [50]. Similar to De Paula et al. [29], the Laissez-faire work proposes that access-control needs to be less restrictive, fit naturally into the end-user’s workflow, and better match current behavior where access control is more continuous and less about definitive *allow* and *deny*.

To help users specify privacy policies in natural language, Brodie, Karat, and Karat built and tested the SPARCLE natural language policy management interface [19]. SPARCLE assists knowledge workers in writing machine-readable natural language privacy policy rules in a guided environment [53]. Vaniea et al. explored the use of syntax highlighting in the SPARCLE interface. They found that when writing rules in natural language, end users need the interface to expressly support the planning/translating and revising tasks normally associated with natural language writing [101]. Using the SPARCLE system, Reeder et al. identified five general usability challenges that policy-authoring systems must address to be considered usable. These challenges include: 1) making default rules clear, 2) communicating and enforcing rule structure, and 3) preventing rule conflicts [78].

Maxion and Reeder observed that, when interacting with access-control permissions, end users rarely care about the individual rules and instead want to see the effective permissions [66]. Effective permissions are the result of considering all relevant access-control rules together to determine whether access will be granted or denied. Maxion and Reeder designed the Salmon system, which showed users the effective permissions and how those permissions were computed. They found that users who used the Salmon system were better able to perform basic policy-management tasks, such as giving a person access to a file [66].

Reeder et al. introduced an interface paradigm for access-control policy management which they call the Expandable Grid. It gives users both a high-level view of all the effective permissions in a system and the ability to drill down and examine any particular permission [80]. They found that end users using the Expandable Grid to perform basic policy-management tasks, such as give Bob access to fileA.txt, were faster and more accurate than users who used the default Windows interface for policy management [80]. Further exploration by Reeder et al. found that the conflict-resolution strategy used by the system to compute effective permissions had a significant effect on end users' ability to accurately make policy changes [12].

The Grey system [8, 9, 10], similar to Beaufour and Bonnet's proposed personal servers with digital keys system [13], is an implemented distributed discretionary access-control system that was constructed and studied in a live environment. The Grey system enables end users to manage access control in a discretionary way while maintaining detailed logs. Every access to a resource requires a certificate-based proof that access should be allowed, thereby ensuring that the logs contain both the access attempt itself and details about why the access was allowed. The system is distributed in that the certificates and proof statements are all developed and stored on smart phones, so no central server is required. The mobility of the devices enables end users to make and change their policy from anywhere, with little effort. A within-subjects study of Grey users found that they created more restrictive access-control policies using Grey than with the physical key system they had used previously. The study also found that Grey users were more able to easily change their policy, which resulted in them giving out less access "Just in case." However, one of the issues with such a system is that more than one person can change the implemented access-control policy without necessarily informing other people who have access. This observation was one of the motivations for this thesis.

Proximity information displays are partially intended to provide additional support

for *ex-ante* control that is not provided by existing technologies. Unlike existing policy-management solutions, proximity information displays will provide end users with a passive way to review their existing access-control policy without having to proactively locate a policy-management interface. By providing end users with information about who could access their resources, I will provide them with an easy way to engage in *ex-ante* control. To my knowledge, there is no existing work which examines placing policy information on the interface to encourage users to engage in *ex-ante* control.

2.3.2 *In-medias-res* control

In-medias-res control, otherwise known as *uno-tempore* control by Stevens and Wulf [94] and reactive control by Bauer et al. [9, 10], is somewhat less studied. Stevens and Wulf [94] define *uno-tempore* control as “The permission is defined at the moment of the access attempt.” Bauer et al. describe reactive policy creation as any policy decision made in reaction to an access attempt or request [10].

In-medias-res control is performed on a case-by-case basis for a specific access in a specific context. Unlike users engaging in *ex-ante* control, a resource owner participating in *in-medias-res* control has an understanding of the context under which the access is taking place and potentially even knows the reputed purpose of the access [9, 48].

In *in-medias-res* control, there is little to no automation on the part of the system. An access request is not approved by the system; instead it is manually or automatically forwarded to one or more users who decide the outcome that the system enforces. Alternatively, a request could be created out-of-band which the resource owner responds to by creating permanent or temporary permissions. As mentioned earlier, Bauer et al. created a physical access-control system called Grey. This system also allows end users to directly request access to offices from office owners who can choose to either allow or deny the request [10]. They found that end users made use of this functionality to manage offices that are accessed only occasionally by people other than the occupant. Mazurek et al. also explored having end users approve or deny access to their files in real time [69]. They found that users responded differently when asked to describe their policy *ex-ante* than when they were asked at the time of the access request (*in-medias-res*).

Another type of *in-medias-res* control is the creation of temporary permissions that can be used only once or for a short time period. Whalen et al. observed a need for granting temporary access to files [106]. Bauer et al. also observed people using *in-medias-res* control to give others one-time or temporary access to an office [9].

Proximity information displays are not intended to support *in-medias-res* control. *In-medias-res* control requires that the resource owner be notified in a timely manner. Because proximity information displays are spatially located on the interface near the resources they refer to, there is no guarantee users will be looking at them at the time when *in-medias-res* control needs to be performed. This makes proximity information displays an inappropriate medium to encourage *in-medias-res* control. However, proximity information displays can help end users engaged in *in-medias-res* control by making it easier to locate the policy-modification interface to make changes. Existing research shows that when trying to solve a problem users tend to start at the problem source, the resource, and iteratively search

outward for a way to solve it [71, 109].

2.3.3 *Ex-post* control

Ex-post control is defined by Stevens and Wulf as “Permissions are checked after access was granted” [94]. In *ex-post* control, the resource owner sets little to no access-control restrictions *ex-ante* and instead relies on accountability to ensure that resources are used in a responsible way. The system logs the details of each access. The resource owner then reviews the accesses after they have happened.

Ex-post control has several major advantages. If the resource owner is in an environment where the majority of users are trusted, managing individual permissions may take more effort than it is worth. As Zhao and Johnson observe “rigid access control delays an organization’s response to the changing markets, resulting in missed opportunities or degraded service quality” [110]. Engaging in *ex-post* control allows the resource owner to let other users use their good judgment and quickly gain access when access is needed. *Ex-post* control also allows the resource owner to evaluate the appropriateness of an access once all the facts are known. As Blakely suggests, “make users ask forgiveness, not permission” [18].

Similar to *ex-ante* control, in *ex-post* control the system is responsible for automatically granting access based on a previously expressed set of preferences. The difference is that in *ex-post* control, the resource owner takes an optimistic view and gives access to all people who might ever need access. The system is responsible for automatically enforcing this policy and the resource owner is responsible for manually reviewing the appropriateness of each previously allowed access.

Ex-post control is based on the observations that end users do not always know who should or should not have access to which resources in the future and that end users have limited time to manually approve and deny every request. In their work Jaeger, Edwards, and Zhang look at the permission-assignment state of individual users in terms of actions that are expressly allowed and actions that are expressly denied. They found that often a significant portion of the access-control space has neither an express allow nor an express deny defined [47]. Rissianen, Sadighi, and Sergot took this observation one step further and applied the idea of access-control spaces to policy creation and enforcement. In their work they argue that unanticipated and unenforceable policy should be enforced with an “Allowed - with override” policy that is enforced via *ex-post* control [82]. Stiernerling and Wulf expanded on this idea by building negotiation functionality into their groupware document-sharing tool. The tool notified users when specific documents were changed and gave the users a technological medium for negotiation and resolution if the change was unacceptable to someone [95].

Stiernerling and Wulf observe that in multi-user collaborative environments, users have need for more complex policy controls than a simple allow or deny. They observe that in collaborative work environments, people have to access each other’s files while at the same time respecting the other person’s privacy. In their work they look at how people use awareness, trusted third parties, and negotiation to handle situations that are unforeseeable or simply outside the abilities of the access-control system to specify. They then add a

tool to their groupware software that allows users to create complex conflict-negotiation rules. The negotiation system allows another user to override the existing access-control permissions provided several requirements, such as an email to the owner or n number of people agreeing to the override, are met [95].

Polvey introduces the concept of “optimistic security” in which the resource owner places few, if any, policy restrictions on the resource and instead relies on accountability and the ability to roll back the system to ensure integrity of the data. In optimistic security the potential accessors are considered somewhat trustworthy, and it is assumed that the majority of accesses will be “good.” If one of the users performs unacceptable accesses, the resource owner has options for recourse via system roll-back and change accreditation. So, while another user can freely read and make changes to resources, the resource owner can easily attribute each change to the person who made it and they can easily return the system to a prior state [76].

Gutierrez et al. proposes a system where end users can negotiate the amount of tracking information visible to the content owners of pages users visit [43]. Content owners would set explicitly the level of tracking detail required to view each piece of content. In this case: detailed information, anonymous information, or no information collected. Users also explicitly state the level of collected log information they find acceptable to be visible to content owners. Each user will be shown only content that matches both their and the content owner’s preferences. Users who are more willing to give up privacy can see more content, and owners who are more willing to display content without tracking will display that content to a wider audience. Though built, this system was never tested with end users.

Proximity information displays are intended to support *ex-post* control by providing information about who has been using which resources. The information allows the resource owner to casually perform an *ex-post* review of the accesses and determine if anything unacceptable is going on without having to proactively open a dedicated interface. There has been limited research on how to construct interfaces that support *ex-post* control and the majority of that research has looked at multi-policy author environments.

2.4 Access-control policy tactics

Both administrators and end users make use of a variety of technologies to create the security effect they want. The tactics they use have been explored in our own work [9, 10, 68] as well as by other researchers [29, 30, 52, 94].

In our own research we conducted a field study of a smartphone-based access-control system in a university environment. We collected usage data of the system and interviewed users every few weeks about their use of the system [9]. We also interviewed them about their ideal and implemented access-control policies [10]. In another study we interviewed home computer users about their access-control strategies for their electronic files as well as paper files in their homes [68].

Access-control technology designs frequently assume that users want to divide the world into two groups of people: those that should be able to perform a specific action on an

object, and those that should not be able to perform the action on the object. However, most users manage their personal access-control using more fine grained distinctions.

We found that when managing access-control, people use a wide range of tactics and social pressure to enact security policies that would not be technically feasible using only system settings alone. The tactics used fall into five main categories: planning for the unexpected, in-the-moment, witnesses, obfuscation, and audit.

Planning for the unexpected – People would give physical keys to another person with the explicit instruction that the key was not to be used except in an emergency. While the other person was trusted, the goal of the permission granting was not to give them daily access. A combination of trust and social pressure was used to make sure the access was not abused. An “emergency” was defined as any unexpected event where the access granter either was unavailable or had remotely authorized the access.

In-the-moment – Privacy and security are often highly contextual; giving access is not only about *what* and *who* but also *why*. People who normally did not want anyone entering their offices would mention several highly context-dependent specific situations where they wanted to allow someone into the office just once for a specific purpose. Those who needed to give in-the-moment access would typically call someone who had an emergency key to open the door or verbally state a key code.

Witnesses – Offices, homes and even folders are spaces that can contain many types of content. Giving a marginally trusted person access to one, even for a specific purpose, was perceived as risky. When giving in-the-moment type access to an untrusted person, the permission granter would require that a trusted person be present. This trusted witness would provide access credentials on behalf of the untrusted person and be physically present to witness all actions that were taken in the space.

Obfuscation – Physically or digitally hiding an object that needed to be protected was a simple low-tech tactic. Hiding required limited understanding of how the security system worked and participants had confidence that no one would target them sufficiently to find the hidden item. Hiding allowed someone to give access in-the-moment without using a third party by verbally telling the accessor where the object or credential was hidden.

Audit – Another tactic was to place trust in a group of people, give them access, trust them to behave correctly, but have a way to audit their actions later. This was enforced either with logs or by placing the object being accessed in an open space visible by many people. This tactic uses minimal technological mechanisms to force correct behavior and instead uses social pressure and the threat of punishment to encourage correct behavior. It also allowed the person whose item it was to make judgments with an understanding of the actual consequences of the actions.

The act of controlling access is not just about allowing someone into the office or not. Issues such as context, purpose, levels of trust, and the ability to reserve judgment until

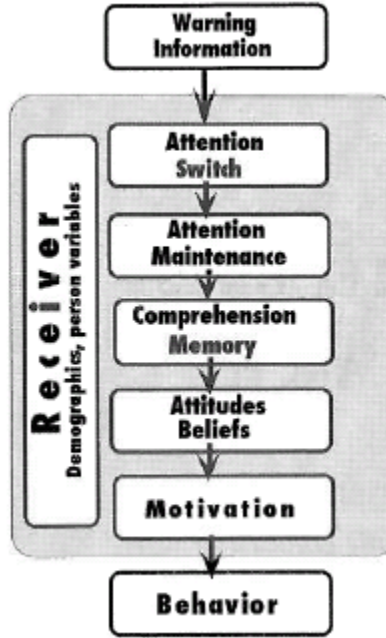


Figure 2.1: Communication-Human Information Processing Model (C-HIP).

effects were known were all major factors. An effective system should support users in these behaviors.

2.5 Behavioral models

Research from cognitive psychology, behavioral economics and the warning sciences provide useful information on how humans react to and think about different situations. In this thesis we design a passive interface component, which is similar to a warning in that it is intended to help people quickly identify situations where they might be in danger. The models presented in this section describe how users process and think about warning interfaces. In this section I will talk about the Communication-Human Information Processing Model (C-HIP) [108] which describes how humans process warning information. I will then discuss an expansion of the C-HIP model called Human In The Loop Framework (HITL) [26] which adapts many of the central principles of C-HIP to the computer security domain.

2.5.1 C-HIP model

Wolgalter proposed the Communication-Human Information Processing Model (C-HIP), pictured in Figure 2.1, for structuring and discussing research about warnings. The C-HIP model is useful for understanding how people process presented information in terms of noticing it, understanding it, deciding if it is important, and finally doing something about

it. According to the C-HIP model, interaction begins with the display of a warning through a channel to the end user. Once the warning has been delivered via the channel, there are several stages users can go through. Each of these stages is described below.

Attention Switch In the initial stage, the warning needs to get users' attention by getting them to look at the warning. To do this, the warning must be designed to be noticeable. It also needs to be positioned such that it can be noticed by users.

Attention Maintenance Once users have switched attention to the warning, their attention needs to be held long enough that they acquire the information presented by the warning. Legibility and form factor can have a strong influence on attention maintenance. If the warning looks difficult to read or unclear users may not dwell on it long enough to attain the needed information.

Comprehension and Memory Even if users look at the warning long enough, they may still not be able to internalize the information from it if they are unable to comprehend it or if it fails to activate relevant information from memory. For example, a "Warning! May contain musca domesstica" sign is useless to someone who does not know that musca domesstica is the scientific name for the common house fly.

Attitudes and Beliefs A fully comprehended warning may still fail in its purpose if it fails to adequately influence users' hazard-related attitudes and beliefs. Beliefs and attitudes form users' current mental frame-of-reference based on users' experiences. For example, a "your files are visible to all people on this computer" warning may be ignored by someone who believes that no one would ever go looking for their files.

Motivation In the final stage, users are either energized to engage in behavior appropriate to the warning or they are not. A motivated user will move on from this stage to the *Behavior* stage where they engage in a behavior appropriate to the warning.

2.5.2 HITL framework

The Human In The Loop (HTL) Framework proposed by Cranor [26] expands and adapts the C-HIP model to the domain of computer security. The work also postulates that through the HITL framework, lessons from C-HIP are applicable to additional communication mediums including notices, status indicators, training communications, and policy communications.

Cranor also introduces the concepts of *knowledge retention* and *knowledge transfer*. If users learn about a particular hazard through a communication are they likely to remember that the hazard exists in the future when they encounter it again? If the same user encounters a similar hazard, are they able to apply the lesson they learned from the warning in a new domain? For example, if a person uses proximity information displays and learns that all of ProjectA is world readable, will they remember that fact later when they try to save secretFile.txt to ProjectA?

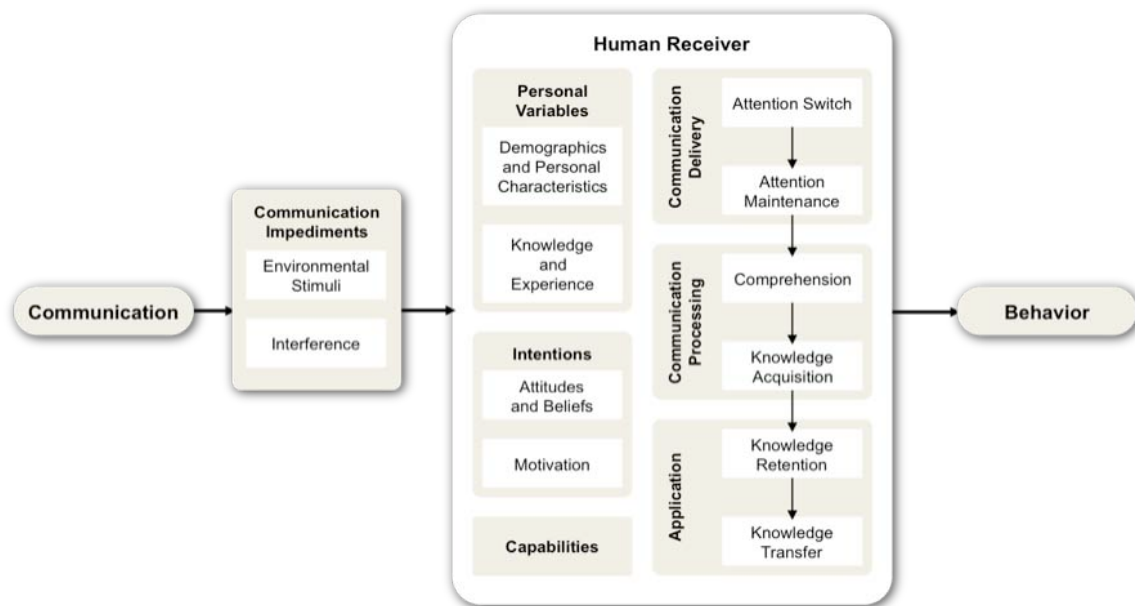


Figure 2.2: Human In The Loop Framework

Chapter 3

Supporting end-user permission management

In related work (Section 2.2), we showed that people frequently have mismatches between the access-control policies enacted on their systems and the policies they want to enact. These mismatches are dangerous to organizations, who might lose valuable data, and end users, who might experience embarrassment or loss. What causes people to have mismatches between what they want and what the access-control settings on their systems are specify? How can we help people decrease the number of mismatches? In this chapter we discuss these two questions and why we ultimately decided to investigate proximity information displays.

We begin our discussion of how people manage access control by defining two types of policies users can have: an *implemented* policy is the current state of all access-control-related settings on a system, and an *ideal policy* is the policy that users want to have enacted on the system. In other words, the ideal policy is what users would like to have happen and the implemented policy is what the system is currently set to do. Any difference between implemented and ideal policies is a *policy error*.

In addition to implemented and ideal policies, we also want to talk about the *understood implemented policy*—what users believe the implemented policy contains. The users' understood policy and the actual implemented policy are not necessarily the same. Users will make decisions based on their understanding of the policy, not necessarily based on the actual implemented policy.

The circumstances which give rise to a policy error can be divided into the following categories. First, users are aware that the error exists but is unwilling to correct it. If users decide that they do not want to correct an error in implemented policy, arguably they have corrected the error by altering their ideal policy; they have changed what they want the system to do rather than changing the system's settings. Second, users are aware that the error exists but are unable to correct it. Third, users may not be aware that an error exists. We assume that users have an accurate understanding of their own ideal policy and can easily compare it to the understood implemented policy. Consequently, the only reason users are unaware of the existence of an error is if their understanding of the implemented policy is inaccurate, such that they inaccurately believe that their ideal and

implemented policies have no conflicts.

If we think about the problem in terms of ideal, implemented, and understood implemented policies, it becomes clear that we need to research solutions that help users realize that they have policy errors and follow through with correcting them.

3.1 Potential research directions

We could have explored many research directions, ranging from removing decisions from users through better prediction technology to developing better warning notices to inform users of likely errors. In this section we present an overview of these possible research directions.

3.1.1 Remove users from the loop

One direction would be to limit users' involvement in policy creation and maintenance by creating better defaults and building systems that accurately predict the users' needs. This approach can never be 100% accurate because only users have a perfect understanding of their ideal policy. However, it could decrease the number of policy errors or flag parts of the implemented policy that users would need to review.

3.1.2 Pull: Wait for user's request

Another direction is to provide information to users only when users requests it. This is how the implemented policy is traditionally presented to users, but there is still significant research that needs to be done in this space. If the implemented policy were presented to users in a way that was easy or even pleasant to interact with, users might proactively choose to look at the implemented policy.

This is a large research space and includes approaches such as building better visualizations of the implemented policy [80], building training materials to help people understand how companies use data [59], visualizing the privacy policies of companies [54], and building interfaces that make it easy to change the implemented policy [66].

One of the problems with this direction is that users often treat security as a secondary task [25], and if they believe that no error exists, they may never seek out information about the implemented policy.

3.1.3 Push: Proactively show information

Another direction is to push information about the implemented policy at users even when users have not asked for it. Because people often treat security as a secondary task, it is unlikely that they will seek out implemented policy information unless they see a clear need to do so. If information is provided to users unsolicited, then users may be more likely to notice and engage with the information.

One of the problems with this direction is the potential for user annoyance or habituation. If users are shown too many pop-ups or emails with security information, they may either become annoyed or start habitually ignoring the information. We already see this behavior with SSL security warnings [91, 96].

There are several ways to push information at users. The user interface of the system on which the policy is implemented can be used to push information at users; alternatively, information could be pushed at users through a channel outside the system.

System user interface One way we could push information to users is through the user interface of the system itself by adding new user-interface elements or displays. These types of displays can be *active* or *passive*. Active displays interrupt users' tasks and force or strongly encourage users to interact with them before progressing. Passive displays present information to users but do not force users to notice or engage with the display. In addition to being purely active or purely passive, displays can occupy various points in the spectrum.

Actively pushing implemented policy information at users is not common. Some mobile websites, such as The Home Depot's mobile website¹, actively ask for access to users' current location via a pop-up. We are not aware, however, of any system that tries to communicate the whole implemented policy in this way.

Using the system interface to passively push information at the users is a method currently used by Facebook, Google+, Flickr, and other online content sharing sites. These sites use interface components to indicate the current state of the implemented policy. However, we are not aware of any research done that specifically studied the effectiveness of these interface components in helping users better understand their implemented policy.

Notification via alternative channels Another way to push information at users is by sending them periodic notifications via a communication channel other than the system user interface. This notification channel could be anything from paper mail to SMS messages.

There is a wide range of ways to notify users about the implemented policy without using system interfaces. Issues such as notification frequency, size of the notification, and layout are all major factors that would have to be investigated to determine best options for a particular system, or even in general.

An example of this kind of approach is banks sending out privacy policies to customers every year. Similarly to helping users gain an accurate understanding of the implemented policy, bank privacy policy notifications are intended to help the customer understand how the bank uses the customer's private data.

3.2 Chosen direction

Each direction we considered in Section 3.1.3 has advantages and disadvantages, but they are impossible to fully quantify without significant research. Designing and testing inter-

¹<http://m.homedepot.com>

faces for every possible direction, while ideal, is not feasible to do in a single thesis.

We decided to investigate pushing the implemented policy, or information about how the implemented policy has been used, via the system user interface. We decided to start with this design approach because it seemed the most likely to assist users without overly annoying them. Testing this design also initially appeared to be easier because we could evaluate it in a single session, while some of the other possible solutions needed to be evaluated over time.

Finally, the different solutions to helping users find and correct policy errors are likely to share many similarities, such as the type of information provided and the way the information is visually represented. Designing and evaluating solutions that push implemented policy information via the system user interface will help us understand how people react to this type of information, which will be useful to anyone trying to test a different but related method.

3.3 Proximity displays

There are many ways to push implemented policy and usage data at users through a system user interface. Based on prior research and our own experiences, we made several decisions that narrowed the design space we would explore.

We were concerned that pushing information at users too aggressively would start annoying them. Active displays block users from accomplishing their primary task until users have interacted with the display. Because computer systems cannot reliably detect policy errors, we were concerned that using active displays would frequently interrupt users when no error was present. We decided to focus our research on passive displays, which show information without actively attempting to draw users' attention.

We also decided that spatial proximity was an important feature. Displaying the full implemented policy on a single active display would take up a large amount of screen space and likely distract users. Instead, we decided to show it in smaller, more manageable pieces. We also wanted to leverage context when showing the interface. We felt that showing people implemented policy and audit information about items they were currently thinking about would improve both their ability to identify errors and improve their long-term ability to accurately recall their implemented policy. Finally, spatial proximity is a known method for assisting users in linking or grouping multiple visual objects [35]. Because of the emphasis on spatial co-location, we refer to the user-interface element where the information is shown as a *proximity display*.

3.3.1 Scenarios

To help the reader better understand how we envision proximity displays will fit into the normal work flow of users, we present three hypothetical scenarios where users interact with a proximity display. These scenarios are intended to demonstrate three use cases for proximity displays: a positive experience where no policy error exists, a scenario where an

error is discovered in the implemented policy, and a scenario where the understanding of the implemented policy was wrong and the ideal policy was adjusted to correct the error.

We expect that security is a secondary task for users, so interacting with the proximity display is also secondary. Each of the following scenarios therefore describes users' context and primary task, showing how users interact with the proximity display as part of their normal work flow.

Alice wants to select some photos for her screen saver. She goes online to her photo albums, and starts looking through her albums including some photos from a trip to Chicago with her good friend Sue. While looking at the album, Alice notices that Sue has recently viewed the album. Alice has not talked to Sue since they returned from the trip. Interested to hear from her friend, she sends Sue an email to catch up on events since the trip.

Joe goes to his online-photo sharing website to share some pictures from his latest vacation. He uploads all the photos into a new album and starts going through them to make sure they are all correctly oriented, have good titles, and generally look good. While going through the photos he notices, via the proximity display, that his ex-girlfriend can see his new photos. Joe is very upset by this and immediately wants to make changes to his privacy policy. He uses the link on the proximity display to open the permission-modification interface for this album. He changes his policy so his ex-girlfriend cannot see his new vacation photos, and then returns to his new album. He uses the proximity display there to double check that the ex can no longer see these photos. He then returns to making sure his new photos are presentable.

Sam likes looking through all the comments people make about her photos. Sam takes great photos and enjoys having other people comment on them. As she is going through her most recently posted album, she wonders with whom she shared this album. Sam generally does not change the access-control settings on any of the websites she uses. She trusts that because the websites are popular, they likely have good defaults. Since she has never had any problem with her online photo albums, she is disinclined to waste time looking through multiple pages of settings. However, she realizes that the settings are visible on the same web page as her photos. She glances at the proximity display, primarily out of curiosity, and realizes that her poetry group can see these photos. Sam did not expect that the poetry group could see her photos, but she decides that she is fine with them viewing her work, so she returns to reading comments. Later she writes a poem about a quirky apple she photographed and points her poetry group to the photo, knowing that they can view it.

3.3.2 Design space

Considering the scenarios above, prior work, and the areas we wanted to assist users with, we decided to limit our design space to information and interface design features that were most likely to help users learn about their implemented policy. The key features and dimensions that we explore are as follows.

1. **Spatial proximity:** As mentioned above, we considered spatial proximity to be important in giving users context and showing them implemented policy information

at a time when they are thinking about the item associated with the policy. We focused on two types of spatial proximity:

- (a) **Place the display on the edges of the screen.** On a sidebar the display is out of the way, but can be easily referenced by users. This spatial placement should limit user annoyance and have minimal negative impact on users' primary task. However, it also has limited co-location; the display could be on the other side of the screen from the content to which it refers.
 - (b) **Place the display adjacent to photo and album thumbnails.** If the display is spatially adjacent, it can be noticed by users who are working on their primary task. This placement has a higher risk of annoying users, but it also places in the display in the closest possible spatial co-location.
2. **Information:** We focused on presenting two different types of information to participants:
- (a) **Display information about who has seen the photos (audit).** In Section 2.4 we observe that people use several tactics to manage security, some of which involve giving other people access and expecting them not to use it. The people given access in this manner are generally trusted, but the person who owns the resource needs to be able to identify whether the resource has been accessed when it should not have been. Presenting information about who has accessed their resource would give resource owners a chance to reassess their policy decisions and take remedial action.
 - (b) **Display information about the implemented policy.** Studies, in addition to our own, have indicated that people frequently have mismatches between their ideal and implemented policies [10, 65]. Displaying the implemented policy to users would improve their understanding of their implemented policy and potentially help them to identify policy errors.
3. **Organization and granularity:** We wanted to design a user interface that supported users in identifying policy errors by better understanding their implemented policies. To this end, we decided that the following three user interface features:
- (a) **Enable layered data exploration by moving some details to a secondary interface.** Proximity displays should help people realize that an access-control problem exists, but they may not be the best mechanism for supporting users in identifying the scope of the problem or correcting the issue. We expect that some information needs to be immediately visible to users, while other information should be *layered* – shown to users only when they interact with the display. Layered data can be anything from tool tips that appear only when users' mouse rests on a component to in-depth permission information explorable in detail on a secondary page.
 - (b) **Display who, not just what group, saw or could see information.** When talking to users and security experts, we noticed people bringing up examples of problems that had occurred because the membership of a group or folder was

not quite what was expected. This was especially true with user groups where multiple people could alter the membership of the group. We wanted to support users' ability to identify individuals who should not have access, not just groups.

- (c) **Show detailed data, not just a single overview icon.** Icons and other small passive indicators take up a small amount of space. If users know what the icon means it can provide detailed information. Unfortunately, end users often do not know what the icons mean. We wanted our displays to have enough information visible that users could, with high accuracy, determine if there was a problem with the policy or not.

Chapter 4

Focus group: User reactions to proximity security information

Section 3.3.2 limited the design space based on prior work and experience. However, variants of these design ideas are too numerous to thoroughly test. We therefore used a series of focus groups to select and iteratively refine designs for further testing.

This chapter describes the initial designs, focus group evaluation methodology, reactions of focus group participants, and the specific design decisions made based on their reactions.

4.1 Interface designs

In Section 3.3.2 we limited the design space to focus on proximity display designs. We describe two types of information that would be useful to users: audit and implemented policy. We also describe four user interface features that we similarly feel will help people better understand their implemented policy and identify policy errors: spatial proximity, layering information, and varying levels of granularity. When creating interface designs to test using focus groups, we wanted to include multiple variations of these types of information and interface features, so that we could gauge participant response to the combinations.

Spatial proximity:

We varied the spatial location of the displays on each interface. We placed the displays:

- On the edges of the screen.
- Adjacent to photo and album thumbnails.

Information:

We showed people two different types of information:

- Information about who has seen the photos (audit).
- Information about the implemented policy.

Layered information:

We made some information on the interface visible without interaction, and other information required users to click or mouse over the display to read it. Layering was used to progressively show more information in more detail. Interfaces that showed very detailed information on the main display had minimal layering, and interfaces that showed abstract or summarized information showed more information when users interacted with the display. In general we tried to present a consistent amount of information in each interface, and used layering to present information we did not want to show on the main interface.

Granularity:

We varied both the amount and detail of audit and implemented policy information shown on each interface. We wanted to show users the names of people, not just groups, who had accessed or could access an album. We also wanted to show detailed information, such as the amount of time a person spent looking at a photo, as a way to help people better understand how their implemented policy was being used. In order to test the impact of showing specific information, we showed information in detailed, summarized, and abstract forms.

Detailed information specifically named individual people and provided low-level details about them. In the case of audit information, details included when an access happened, how long the album or photo was looked at, and what group membership gave users access. In the case of implemented-policy information, details were limited to group membership and allowed actions.

Summarized information provided a high-level view that typically referred to group and album names and rarely mentioned individual names with specific information. Though high level, this information used specific numbers and names. In the case of audit information this was the number of times the album had been viewed by the different user groups, or by the geographical locations of the users. In the case of implemented-policy information, this was how many people could view the album, how many groups could view the album, or which groups could view the album.

Abstract information tried to convey a sense of the the implemented policy and who visited the site without giving specific numbers. In the case of audit information, the albums could be ordered from most to least viewed, the size of the album thumbnail could reflect the number of views, or the size of the album name in a word cloud could indicate viewing frequency. In the case of implemented-policy information we used green, red, and yellow colors to indicate albums that were public, private, or had custom settings. When further details were not provided, such as which groups had access under custom settings, this information was considered abstract.

The interfaces we decided to use in the focus groups were designed to loosely resemble real interfaces, but left out many of the details an actual interface might have, such as the tools for photo editing or the Facebook and Twitter share buttons. The permission information shown was also mildly exaggerated, taking more screen real estate than would be normally feasible and showing more details than should be needed. We wanted partic-

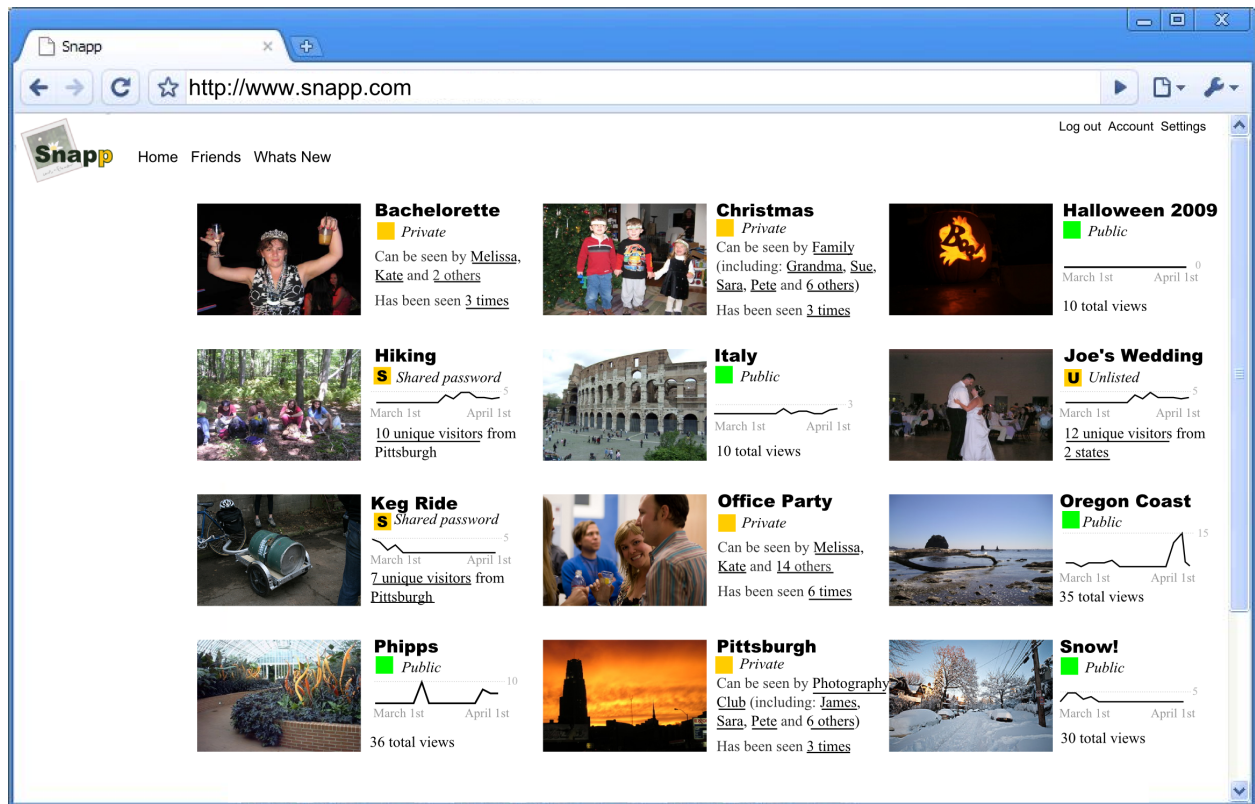


Figure 4.1: Example interface typical of the ones shown to focus group participants. This interface shows detailed implemented-policy information and summarized audit information adjacent to the album thumbnails.

ipants to focus on the permission information we were showing them and have minimal distractions from other user-interface components. We also wanted them to comment on multiple aspects of an interface, with the end goal of decreasing the size and amount of information shown to a more manageable level based on participants' reactions and perceived usefulness. Figure 4.1 shows an example of the type of interface we showed to participants.

4.2 Methodology

The methodology used was typical for research involving focus groups. Participants were asked several initial questions to get them thinking about privacy and security, and then they were asked to comment on each of several website designs. Based on feedback, we altered the designs between focus groups so each group saw a slightly different set of interfaces. We transcribed the notes and the audio and grouped the comments by concept.

4.2.1 Participants

Participants were recruited from an existing pool of people in the Pittsburgh area who had previously indicated interest in behavioral research studies. We recruited a total of 28 participants for five focus groups. Each group had between four and six participants. The majority of participants were students, and all had previously shared photos online.

4.2.2 Protocol

The focus group was conducted in a conference room at Carnegie Mellon. The sessions were audio recorded and lasted for an hour. To get participants thinking about security and privacy, we started by asking participants:

1. About the websites they used to share photos.
2. The last thing they shared online.
3. A time they discovered that someone they did not want to share with could see their content. (Not all participants were expected to answer.)
4. A time they tried to share content and had not been able to due to technical issues. (Not all participants were expected to answer.)

Participants were then handed two pages with cartoons on them (Figure 4.2) that illustrated use cases for providing users with privacy and security information in an easy-to-notice way. Participants were then asked: “Can you imagine an instance where you or a friend might experience a situation like those Alice and Joe encountered?” Participants were encouraged to briefly discuss the answer to this question amongst themselves.

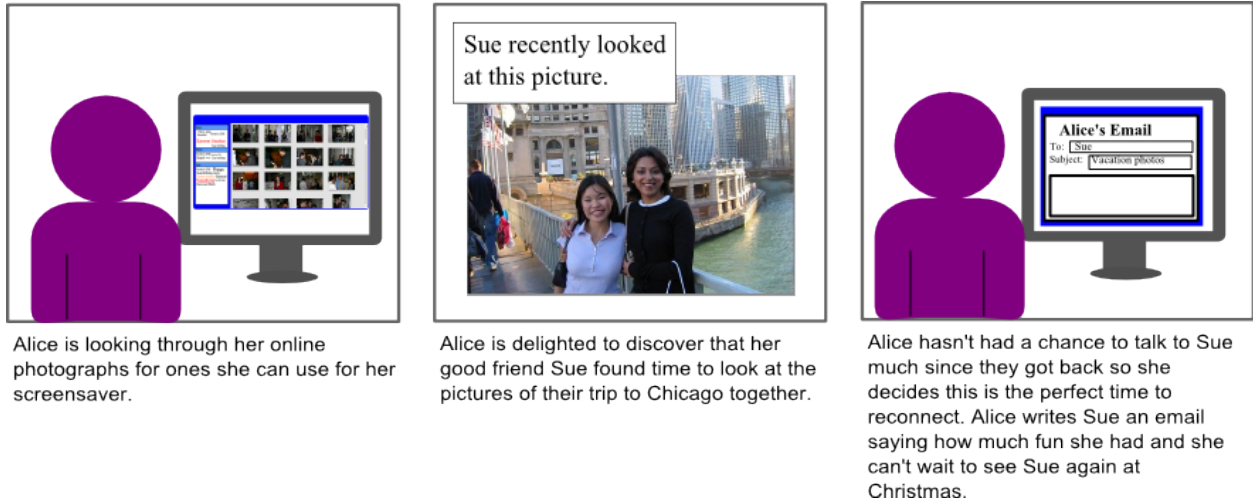
Participants were then given a packet of website designs. Participants were told that they would be going through the packet as a group and were asked to not look ahead. For each webpage in the packet, the researcher gave the participants a brief presentation of the site, its features, and any interactive components. If participants had any questions about how the site worked, they were allowed to ask them at this point. The researcher then asked the participants to fill out the questions included in the packet in silence. When everybody was finished, the researcher started the conversation by asking each participant what the best and worst parts of the website were for them and encouraged them to discuss the answers. Participants were informed that they could write in the packet at any point, so if they were unable to voice an opinion, they were welcome to write it.

At the end of the study session the researcher asked each participant what their favorite and least favorite website design was and why.

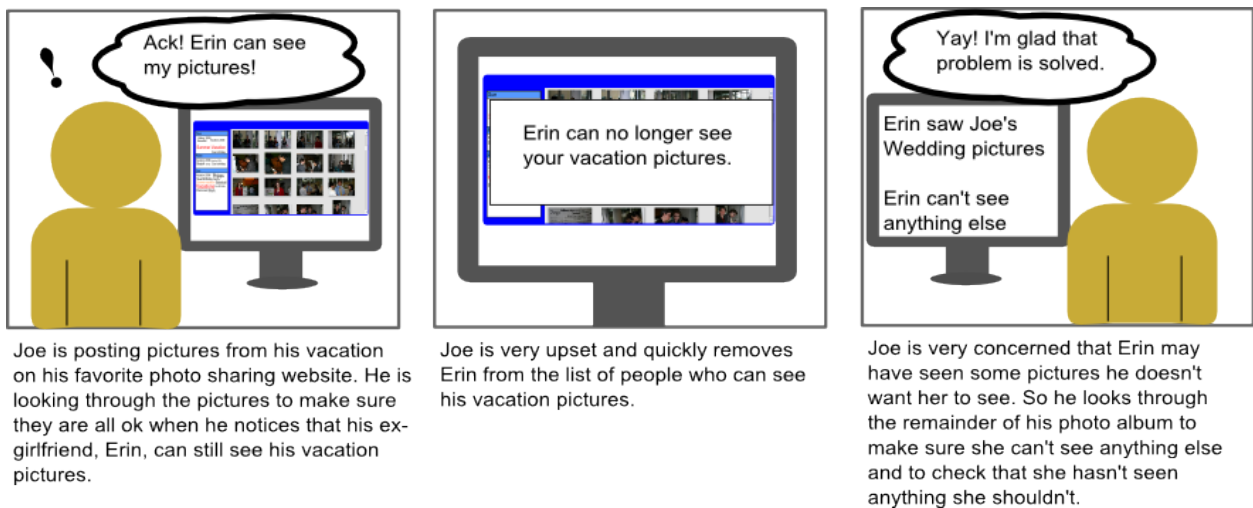
The researcher’s dialog was scripted as much as possible. A detailed script for the first focus group can be seen in Appendix A.1. The remaining focus groups used nearly identical scripts.

4.3 Results

Participants’ comments were collected and transcribed from the packets and the audio recordings. The transcribed comments were sorted by topic using an affinity diagramming



(a) Usage scenario where Alice notices that a friend has viewed one of her photos which results in a positive experience.



(b) Usage scenario where Joe notices that an ex-girlfriend can see his vacation photographs. Joe views this negatively which causes him to update his access-control policy.

Figure 4.2: Usage scenarios illustrating how an end user might use proximity displays both to cause a positive social experience and to notice an issue.

methodology [17]. The high level takeaways from the studies are detailed below.

4.3.1 Why is privacy important to me?

Similar to what others have seen from the college-age demographic, we observed an interesting mix of participants who consider anything posted on the Internet to be potentially public and participants who want fine-grain control over who can see what [2].

The difference between these two viewpoints appeared to be based in two mental models of how information uploaded to websites would be treated by companies. Participants who considered everything on the Internet to be public had the view that the companies they gave their data to were going to lose it, change the privacy policy without warning, or in some way accidentally expose the users' data. These participants did seem to use individual privacy settings, but felt that the settings expressed what they wanted to have happen, not what was actually going to happen. For them it was important to only place data online that would have limited negative impact if it became public, and spending significant time setting up a detailed policy was a waste of time. This viewpoint was more predominant in the younger participants and appeared to be drawn from experiences with companies like Facebook, which has been known to retroactively change the visibility of pre-existing data, a practice which contributed to a ruling against them by the Federal Trade Commission [3]. Existing research also shows that people lose faith in a company after what they consider to be a betraying change to the company's privacy policy [22]. Participants with the other mental model were more optimistic about how online companies would protect their data. These participants generally agreed that putting embarrassing photos online was not smart, but they also felt that actively managing their privacy settings would make a difference and would be honored by the websites.

Participants who consider everything online to be public tended to dislike the idea of proximity information displays because they cluttered the screen with "useless" and "creepy" data. These participants also tended to refer to people who want to control access to each album or picture individually as "control freaks," or "micromanagers." This is very similar to the culture Gaw et al. observed that people who take extra security precautions tend to be considered "paranoid" [38]. These participants tended to respond to privacy and security information by suggesting that we remove it or change it to a small icon similar to what is now used on Facebook and Google+. One such participant wrote in his packet "Want a lock sign."

Participants who felt there was a need to control access to pictures tended to like the idea of locating information about who could see and who has seen the photos in close spatial proximity to their photos. They felt that proximity information displays made the information easier to find and easier to understand. In the words of one such participant: "I like that you can immediately see who viewed your pictures without necessarily accessing the album in question." They expected that being able to see the controls would better enable them to identify issues and be more aware of their settings.

4.3.2 Who has viewed my photos?

One element we wanted to experiment with was the effect of showing people not only who could see their photos, but who had seen their photos. We refer to this information as *audit information*. We felt that this type of information would assist people in re-evaluating their policies and finding issues by giving them a sense of how their permission settings were actually being used. However, our focus group participants strongly disliked being shown detailed audit information about personal photos.

A participant in the fifth focus group explained the concern: “Too much specific information about who has been seeing what, it makes me uncomfortable as a poster and as a viewer. I do not want stalking information available.” Several participants across the focus groups brought up the term “stalking” in reference to detailed audit information. Participants felt exposed and uncomfortable by the data both from the perspective as a viewer, but also as a photo owner. The issue stemmed not only from the data itself but also its location. Participants felt that this data might be acceptable if the photo owner had to go out of their way to find it, but by placing it on the main interface we were encouraging stalking and taking away their right to choose to see the information or not. One participant explained it as “You are forcing me to stalk my friends.”

The more detailed the audit information, the more concerned participants became. One display showed the date, time, and duration of every view. One participant circled this display in their packet and wrote: “Creepy!!!” The concern was not just with the data exposure, it was also with how others might misinterpret the data. A specific example concerned the duration of a view; one participant commented that it was creepy that someone would view the photos for an hour, then a different participant pointed out that they might have walked away from the computer. This sparked a conversation about how data might be misinterpreted and thereby cause someone to think ill of another when no wrong had actually been committed. Participants were very concerned about how others might misinterpret actions taken online.

Participants were also concerned about trying to extract permission setting information from the audit information. They did not want to accidentally confuse situations where a person used to have access, and therefore showed up in the audit information, with the person currently having access. Early focus group participants primarily saw interfaces that showed detailed audit data or they saw detailed permission setting information, but not both. In later studies; we placed specific current permission information with specific audit information and saw less concern about confusion if someone could currently see the album or not.

While the majority of participants did not like detailed audit information about personal photos, there were a few participants who saw the information as potentially valuable if used in the right places. These participants commented on the usefulness of understanding who had seen their photos and the interface components we had added that explained why that person could view the photos. In the words of a participant: “I like the additional features that help you see who viewed your photos and why (what groups they are a part of etc.)”

Participants hated specifics about who had viewed their photos but they loved high-

level statistics, particularly information that helped them answer the question “Is my photo popular?” On interfaces with no audit data shown, we would get comments asking for the number of views or some very high-level sense of what was popular and what was not. Because of the push back about audit data being specific, we tried showing participants a display which graphically sorted items into most to least viewed, and another that showed sparklines. However, participants made it clear that they really wanted the specific number of times their photo or album had been viewed.

4.3.3 Who could see my photos?

Participants were more positive about seeing information concerning who could see their photos than who had seen their photos. They considered this information useful and enabling. Similar to other topics, what participants liked or found useful about this information depended on their mental model of how effectively companies would honor the settings associated with the participants’ data.

Participants who were convinced that all their photos were essentially public anyway generally considered the majority of the information shown to them to be “irrelevant” or “unhelpful.” They especially disliked the amount of screen real estate the various displays required. These participants preferred the idea of using an icon or something very small and high level to express the policy. If they needed to know more, they felt that they could easily click through to some other screen and see it.

Participants who considered permissions to be worth setting liked how easy it was to see their policies. They felt that the displays made it easy to change their policies, indicating that the idea of using displays as a segue to modifying permissions made sense. In the words of one participant: “I like the control over who can see [the pictures] and how simply that control is apparent.” Participants liked the comprehensive approach to policy display. Interfaces that used words to explain the information shown tended to be liked by participants and considered easy to understand. More detailed interfaces or ones that made heavy use of icons were less well liked. Participants were concerned that they, or less computer savvy people, would not understand the meaning.

While the idea of showing who could view a photo or album was generally liked, participants were concerned with the detail and space required to show the information. They disliked showing individual names of people who could view the photos, because they were concerned that the visualization would not scale well. They also talked about how they thought about the people they shared with as groups, not individuals: “I think about my friends in clusters, bicycling, activist, college.” Participants were also concerned with their ability to think about that number of people at once, and the consequences if they forgot someone: “When controlling who can see what on a per person level you have to be aware of every person. If someone is not able to view something you can have hurt feelings even if it is easy to change settings.” Participants were of the opinion that by using groups, the interface would be easier to glance at and the policies easier to manage.

4.3.4 Proximity displays in personal and work environments

In the previous sections, we have discussed participants’ reactions to different aspects of proximity displays in an online personal photo-sharing environment. In the last two focus groups, we added two website designs based on document sharing to the end of the website lineup. We were curious to see whether participants would react differently to proximity displays in a more work-oriented domain. Participants’ comments on the document websites were nearly a complete reversal of their previous comments and concerns.

Detailed audit information, which was heavily disliked in the photo-sharing context, was in high demand in the document sharing context. It was considered creepy and stalkerish to look at who had viewed photos, but it was considered very useful to see who had interacted with a document and the exact type of interaction conducted. If changes had been made to the document, then participants wanted to see the exact changes. Unlike photos, the more specific the audit information, the better.

Detailed information about who could view documents was also in high demand. We showed a version of the Expandable Grid [80] to the participants in focus groups 1, 2 and 3 on the photo-sharing website, and focus groups 4 and 5 saw the grid on the document-sharing website. Participants disliked the grid on the photo-sharing website because it was “too much information,” but loved it on the document-sharing website: “Loved the grid concept, good to know if everyone in group has been looking.” Participants who considered the grid “too much for the main [document] interface” suggested it be moved to the documents main page, not completely removed from view like the photo-sharing participants had suggested.

4.4 Design implications

We look back at our chosen design space in terms of people’s reactions. For each type of information and user interface feature, we briefly discuss how people reacted to the different possible designs. We also discuss how we altered our design decisions to account for participant opinions.

Spatial proximity

Placing information on the main interface was generally liked. There was not a clear difference among the spatial positions of proximity displays that we tried. Occasionally, when information was shown on the sidebar and album thumbnails were shown in the center, participants were uncertain which album the sidebar information represented. However, that was the only real spatial-based confusing point. Considering the lack of a major difference in spatial positioning, we decided to test several positioning options in our later evaluations.

Information

We tested two types of information, which we showed with varying levels of specificity.

Audit information: Detailed audit information about personal photos was considered creepy by participants. However, they were interested in summary level audit information. They were also interested in seeing detailed audit information about documents. Considering participants' opinions, we decided to initially focus on displaying implemented policy information and later evaluated displaying audit information separately from implemented policy information.

Implemented policy information: Participants generally considered implemented policy information to be useful. They thought putting implemented policies on the main screen better enabled them to identify errors and understand how their information would be used. Considering participants' opinions, we decided to focus our designs on displaying implemented policy information.

Organization and granularity

We also tested the potential effectiveness of several user-interface features: spatial proximity, specific names of people/photos, layering, and comprehensibility.

Enable layered data exploration by moving some details to a secondary interface. Participants liked the idea of layering data. They viewed screen real estate as an important resource and did not want to spend too much of it on displaying implemented policy or audit information. These opinions reinforced our existing opinion that proximity displays need to be small and take up minimal screen real estate.

Display who, not just what group, saw or could see information. If information was shown on the interface, users wanted it to be specific. They were fine with detailed and summarized information but did not like abstract information. Additionally, participants preferred names of groups and albums over individual names of people and photos. They felt that groups and albums were useful to them. Considering these opinions, we decided to focus on interfaces that primarily used or emphasized group names over individual names.

Show detailed data, not just a single overview icon. When participants turned the page in their packets and viewed a new web page design, they would sometimes comment on something that stood out for them, such as: "Wow Alice can't see anything." When designing the initial interfaces to evaluate, we put emphasis on keeping sufficient information on the interface so that people can potentially identify issues without having to interact with the interface.

4.5 Conclusion

In this chapter we have explored how people react to different types of access-control related information when it is presented on the main screen of a photo-sharing website. We found that the mental model participants have concerning how their data will be treated and protected by websites impacts the types of information they perceive as useful. Participants who feel that their data will likely be exposed have limited interest in showing permission

information on the main screen, while those who feel that their settings will be honored are more interested in seeing this information. Regardless of their belief in permission setting effectiveness, participants found information about who had previously seen photos to be creepy and similar to stalking. Conversely, information about who could view their photos in the future was considered to be useful and enabling.

Chapter 5

Proximity access-control information displays

In the previous chapter, we evaluated design decisions by presenting focus groups with user interface mockups. Here, we take their feedback into consideration to build a functional interface. The interface is implemented as a plug-in for Gallery 3, an open-source photo sharing system. Over the course of four evaluation studies, discussed in subsequent chapters, we gathered and incorporated additional feedback. This chapter discusses all variants of the design that were tested, including the proximity interfaces, permission-modification interfaces, and changes to the underlying access-control system.

5.1 Design space

Based on the feedback from our focus group participants, we further limited our design space and made specific decisions about the exact interface to implement and evaluate.

Spatial proximity:

The idea of spatial proximity was generally liked by focus group participants, so we wanted to keep displays spatially co-located with the albums and photos they refer to. However, it was not clear if placing information on the sidebar or adjacent would be more effective. We decided to create a single design and evaluate it on the edges of the screen (left, right, and top) and under photo and album thumbnails. These two spatial locations were most promising, and we felt that testing placing the display on other sides of the screen, or to the right of the photo or album thumbnail, would provide us with limited additional data.

Information:

Focus group participants liked seeing detailed implemented policy information, but were concerned about detailed audit information. Because of this distinction, we decided to focus on the implemented policy.

Despite participants' misgivings about detailed audit information in photo sharing domains, we still wanted to test if audit information could be used by participants to identify policy errors. We feel that audit data could be very valuable to people, and participants were more positively inclined towards audit information in other domains, such as document sharing. Therefore, we created an interface which showed audit information to evaluate its potential ability to help users find policy errors.

Organization and granularity:

Enable layered data exploration by moving details to a secondary interface. Focus groups liked the idea of layered data because they wanted the display to remain small. This opinion reinforced our existing view that proximity displays need to be small and take up minimal screen real estate. Information, such as group membership, which required a large amount of screen real estate, was placed on a pop-up display.

Display who, not just what group, saw or could see information. Focus group participants preferred names of groups and albums over individual names of people and photos. They felt that groups and albums were useful to them. When creating the implemented-policy display design we decided to show group names on the main display, and only provide individual names on a secondary interface (layered). For the audit-display design we emphasized the group names and showed a small number of individual names, with the remaining individual names visible on mouse over (layered).

Show detailed data, not just a single overview icon. Focus groups appreciated having sufficient information on the proximity display to identify potential issues. When creating the display design to test, we made sure to represent the details of the implemented policy and not overgeneralize any important features. (We also tested a condition with a single icon as a comparison.)

In the following sections we describe the designs we implemented based on the decisions described above. While some of our following design decisions are mildly constrained by the platform we selected, we endeavored to keep these decisions and the opinions of focus-group participants in mind while designing.

5.2 Platform

To test our proximity display design, we decided to implement it as a plug-in to Gallery 3 [1]. We chose Gallery 3 because it is an open-source photo sharing system, so we were able to modify how access control was handled. Gallery 3, and its predecessor Gallery2, are primarily used by people with technological backgrounds, and is not a system the general Internet population would be familiar with. Additionally, Gallery 3 was released in October of 2010 and the user interface was a significant departure from Gallery2. Since this research was started in February of 2011, we were confident that the majority of users had no prior experience with the Gallery 3 user interface. Comments from users during our





studies supported this assumption; no participant mentioned prior experience and most users asked if we had designed Gallery 3.

5.3 Proximity display plug-in design

The proximity-display plug-in interfaces used in this thesis fall into four categories. The first two displayed implemented policy information in a grid-based design and a list-based design. The grid-based design showing implemented-policy information was shown to users as part of the pre-study and eye-tracker evaluation studies. The list-based design showing implemented-policy information was shown to users as part of the lab and online studies. We moved from grid to list designs to reduce user confusion over the symbols and to make the display more compact. The third was a list-based design showing information about who had previously interacted with the photos (audit). This design was pilot tested during the pre-study, eye tracker, and lab studies. It was formally evaluated only in the online study. The fourth design was an emulation of Facebook’s icons. This design was intended to test the effectiveness of a proximity-display design used by a popular company. The Facebook-icon display was only tested in the online study.

5.3.1 Implemented policy shown in a grid-based design

The idea of using close spatial proximity to link concepts is well known and part of Gestalt’s principles [35]. These principles describe how humans visually group and associate objects; visual objects which are in close spatial proximity are considered to be related. We use this principle here to bring access-control information into the immediate context of the user’s work-flow. We want checking and changing access-control settings to be as natural as checking other spatially linked features such as titles.

The initial proximity display design (Figure 5.1(a)) shows access-control policy in grid form, with each row of the grid showing the permissions a particular group has to the album in question. Mousing over a group name will reveal the group members, and the permissions are indicated by icons (view , edit , and add photo ). Greyed-out or missing icons indicate lack of permission; icons with a yellow dot indicate that subalbums or photos do not have consistent permissions (e.g., the group may have a specific permission on some subalbums but not on others). If a group cannot view an album, then all other permissions are also unavailable. Figure 5.1(a) is an example of such a display taken from our under-photo condition. Mousing over any icon on the proximity display results in a tool tip with an explanation of the permission in its current context. In Figure 5.1(a), for example, mousing over the  icon next to “Friends” would display “The group Friends cannot view Animal Shelter Shared Albums.”

The grid design is based on work by Reeder et al., who successfully used a combination of grids and effective permissions (discussed below) to make it easier for users to manage file permission settings [80, 81]. Participants were able to use his grid design to get a quick sense of permission settings and focus on important components easily.

The decision to use icons in the grid is based on work by Tam et al., who tested

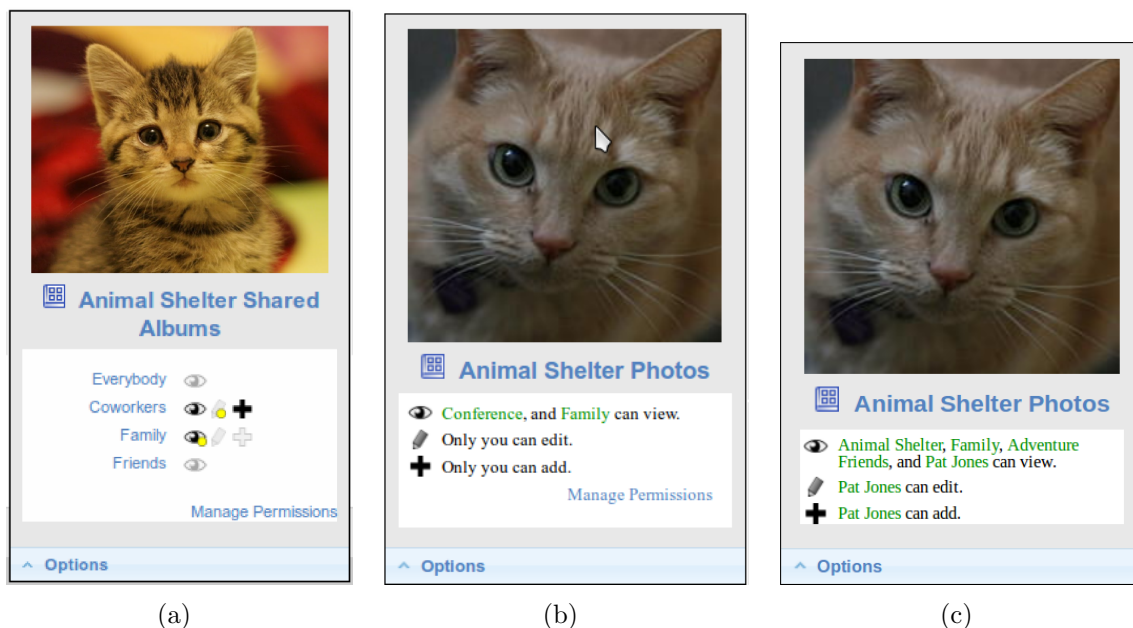


Figure 5.1: The proximity displays shown to users in the four evaluation studies. The display used in the first and second studies (a), is based on a grid-style design. The displays used in the third (b) and fourth (c) use a list-based design. Displays (a) and (b) include a “Manage Permissions” link; participants were rarely observed to use the link, so it was removed in design (c).

multiple privacy notification layouts, intended to be shown during application installation, against participant comprehension speed [97]. They found that layouts that used visual icons allowed users to find data quicker and were preferred by the users. They also found that participants performed better when permissions were organized by action icons. Tam’s work used a different set of actions and the layouts were not for proximity displays, however, we have a similar goal in that we want people to comprehend our displays quickly.

5.3.2 Implemented policy shown in a list-based design

The list-based design came out of our experiences with the grid design. We found that while the symbols readily made sense to most people, other people became confused and therefore set the policy incorrectly. The displays which explained each icon when it was moused over proved to be of limited assistance to participants with incorrect mental models. The researcher also observed that trying to comprehend the implemented policy by quickly glancing at the grid was challenging. The grid design aesthetic meant that users had to first read the group name on the left and then mentally connect it to the icons next to it. While this was easy to understand, we felt the design was overly challenging for users to take in “at a glance.” Too much focus was required to parse the information. Finally, the grid based display was fairly large and had limited scalability.

The list-based design (Figure 5.1(b)) incorporated the same icons as the grid (view ,

edit ✎, and add photo ➕). However, these icons were displayed statically and were only intended to graphically indicate the action the list refers to. These choices were consistent with layouts Tam et al. showed to be effective [97]. Even if no groups had access, the icons were shown. If child albums had different permissions from the parent, this was indicated on a completely separate line that read: *some subalbums have different permissions*.

The list itself is designed to be both understandable and scalable. Where the grid design was organized with each line corresponding to a group, this design made each line correspond with an action (view, edit, add). The icon at the beginning of each line visually indicates the action and makes it clear where each line begins. The group names are embedded in a sentence that clearly states, in words, what action these groups can engage in. The group names are shown in a different color than the surrounding text to make them easy to visually separate from the static parts of the sentence. This was intended to assist participants in finding group names quickly. The design is also scalable. If there are too many groups to list, then the display shows “and 3 more groups” as a link so participants can easily click and see all groups.

In the online evaluation study, we decided to make the display appear on the screen at all times. In the earlier designs (Figures 5.1(a) and 5.1(b)), when the display was located under an album or photo thumbnail, it was only visible if users placed their mouse over the thumbnail. When it appeared elsewhere on the screen, it was always visible. In the online study, we made the display under album/photo thumbnails always visible to make the displays more comparable with other designs such as icons (further discussed in Section 6.4.1). Researcher observation of how people used the designs in Figures 5.1(a) and 5.1(b) suggested that users were ignoring the “Manage Permissions” link and instead using the link on the Options menu. Log data also supported this conclusion. To save space, we removed the “Manage Permissions” link from Figure 5.1(b), resulting in Figure 5.1(c).

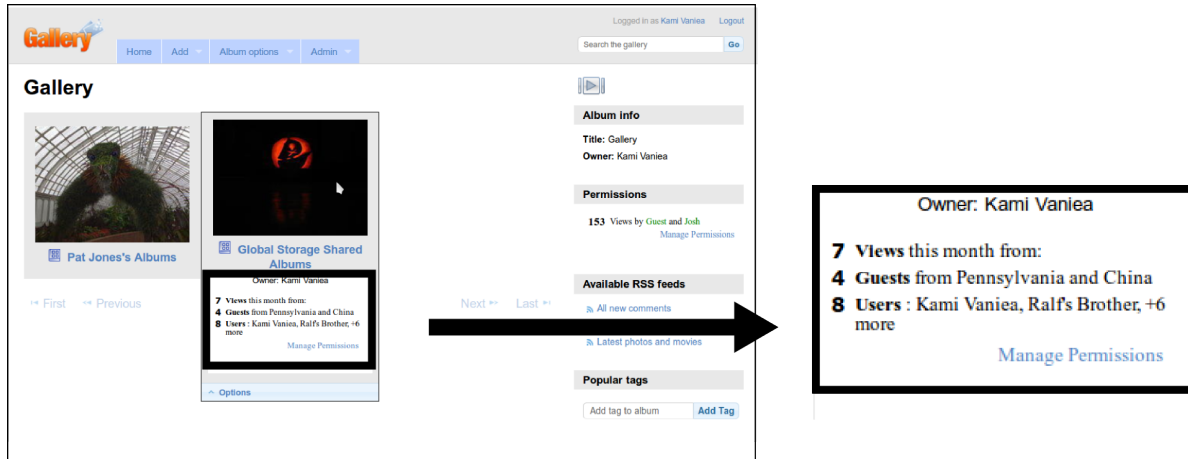
5.3.3 Audit information shown in a list-based design

One of the goals of the thesis is to help people not only be aware of their implemented policy, but also assist them in re-evaluating it based on past performance. We received some negative critiquing from the focus group participants because they felt that the idea of showing detailed audit information was creepy and a bit stalkerish. However, we received more positive feedback when we discussed using it in a work type environment. We feel that this is an interesting direction to explore.

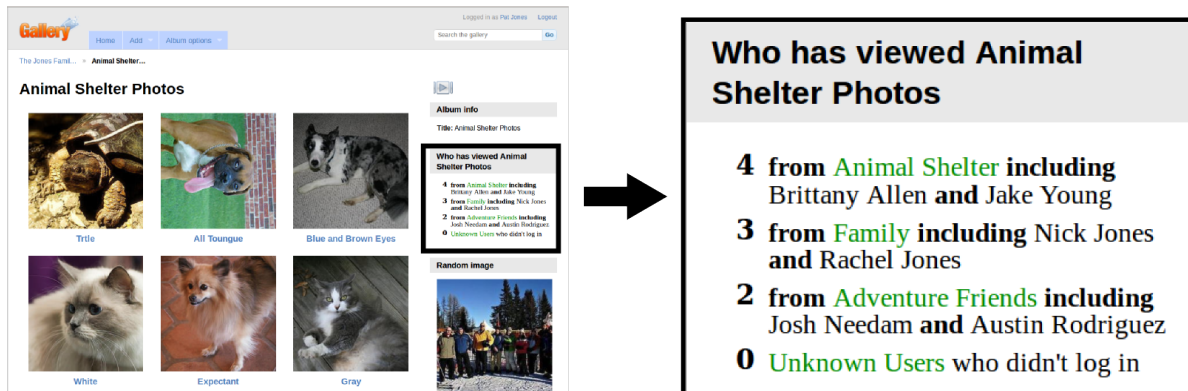
The resulting list-based proximity display design with audit information, shown in Figure 5.2(b), was designed based on the reactions from focus group, pre-study, eye tracker, and lab participants. We also took into account the way people interacted with the grid- and list-based proximity interfaces.

The focus group participants were very clear that displaying detailed information about who had viewed their photos made them feel uncomfortable. However, they were interested in understanding how popular their albums were. They also felt that audit type information was more acceptable if it was either general, or on a separate page where someone would have to actively attempt to find it.

Our initial design, shown in Figure 5.2(a), was intended to make it clear how many



(a) Audit proximity display shown in the 2nd evaluation study.



(b) Audit proximity display shown in the 4th evaluation study.




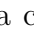
Figure 5.2: The proximity displays showing who had accessed the album (audit). Unlike the displays shown in Figure 5.1, which show who could access the album in the future, these displays show who has accessed the album in the past. Figure (a) was pilot tested during evaluation studies 2 and 3, resulting in the Figure (b) design, which was evaluated in the final evaluation study (study 4).

times the album had been viewed (popularity), and who had seen it. When we tested this design with users, they rarely interacted with it. When asked, they stated that the information shown did not help them find errors, so they had ignored it. Our focus group participants had similarly expressed concern about the helpfulness of displaying individual people’s names, stating that they thought about their friends as groups.

The display used in the online evaluation study (Figure 5.2(b)) is organized by groups with each group displayed on a single line. Only groups whose members have accessed the album are shown, with the exception of “Unknown Users,” which is always visible. When we talked to focus groups, we found that they cared about random people on the Internet viewing their photos. To add some consistency between the displays, and to reassure users, we added “Unknown Users” to every display. Group names are shown in a contrasting color so that users can easily see which groups have access. The number at the front of the line indicates the number of people in that group that have viewed the album. Focus group participants expressed a disinterest in details, including people’s names, but we wanted to encourage scenarios similar to the Alice use case in Section 3.3.1. To balance these conflicting goals, we decided to put the names of the group members who had accessed the album in a non-bold face font so as to de-emphasize them. Additionally, we only list the people who have seen the album on the album page, but not on any of the photo pages.

We observed users working with the audit-information design during the eye tracker and lab studies, and tested it on the online evaluation study. We initially wanted to test the audit display during the eye-tracker study, but when we put it in front of pre-test participants, they ignored it as irrelevant information. This type of interface is challenging to test in a role-play style lab environment because participants are working with a fictitious ideal policy and are only in the lab for 1-1.5 hours. In that time frame, there is not really time to observe participants re-evaluating their permission choices. When we tested the audit interface in the the online evaluation study, we focused only on participants’ ability to identify errors using the interface, and left the question of policy re-evaluation for future work.

5.3.4 Facebook icons

The Facebook-icon display was intended to simulate Facebook’s access-control permission indicators as closely as possible. We decided to use Facebook’s user-interface design because it is both a very popular site for sharing photos and its user interface-design is very different from our own. Facebook uses a set of icons to express the privacy policy associated with albums. An album can be publicly visible () , visible only to the owner () , visible only to friends () , or a custom settings () . Similar to Facebook’s user interface, we placed the relevant icons under each album thumbnail, and when the album was opened we placed the icons in the upper right hand corner. Mousing over the icon resulted in a pop up listing the groups who had the right to view the album. However, clicking on the icon resulted in our permission-modification dialog (Section 5.4.2) rather than Facebook’s drop down menu. Since we are testing if people can notice errors, rather than the impact of the permission-modification interface design, we felt it was more important that the permission-modification interface be consistent across proximity-display designs than for

it to be consistent with Facebook.

5.4 Access-control permission-modification interface

We were concerned that the default permission-modification interface for Gallery 3 might be confusing for end users. We also wanted to make sure our proximity information display design sufficiently matched the permission-modification interface so as to not confuse our users. To this end, we created two permission-modification interfaces: a full-page interface which shows the access-control policy for all albums on a single page, and a dialog interface which shows only the access-control policy for a single album in a pop-up dialog.

5.4.1 Full-page interface

The full-page permission-modification interface (Figure 5.3) displays the access-control policy for all the albums in Gallery 3 on a single page. This interface was designed based on Reeder’s Expandable Grid [80], which was shown to be effective in assisting users in understanding and accurately managing their access-control permissions.

Each row of the grid is associated with an album, and each column is associated with a user group. Each cell contains one or three icons indicating the actions this group can currently perform on this album (black icons), as well as the actions the administrator can grant but are not yet allowed (grey icons). The add and edit actions are not possible when the view action is denied, so the icons for these actions are not shown at all. For example: The *Animal Shelter* group cannot view *Amanda’s Wedding* photos; consequently they also cannot add or edit this album. Giving the *Animal Shelter* group the ability to add without the ability to view would not actually give them any rights, since viewing is necessary to add. Hence, the icons for add and edit are removed.

Users can grant/deny any action by clicking on the icon, which will change it from black to grey or vice versa. To assist users in understanding the icons, the interface includes a legend in the bottom left.

5.4.2 Dialog interface

The dialog permission-modification interface (Figure 5.4) shows the permissions associated with a single album. The display opens as a JavaScript dialog box, so the user can easily view permissions and make changes without having to switch pages.

The design of the dialog is intentionally very similar to the grid-based proximity display. We use the same icons, organization, and mouse-over effects. Similar to the full-sized permission-modification interface, users can change the access-control policy by clicking on any of the icons.

-- back to the ...

Gallery

Dashboard
Settings
Modules
Content
Appearance
Users/Groups
Maintenance

Manage Permissions

[Return to The Jones Family](#)

▼ The Jones Family

▶ Amanda's Wedding
▶ Animal Shelter Photos
▶ Building Jumping
▶ Christmas at Jennifer's
▶ Family Calendar
▶ Grace's Birthday
▶ Jennifer's Baby
▶ Jungle Flight
▶ Safari
▶ Ski Trip
▶ White Water Kayaking

Everybody on the internet	Animal Shelter	Family	Adventure Friends	Pat Jones

Legend

- View album/photograph.
- Edit album/photographs.
- Add new album/photographs.
- Grayed out icons indicate this group does not have this permission on this album.
- Yellow dots indicate that albums within this one have different permissions. Click the album name to see albums within it.

Figure 5.3: Full-page policy-modification interface used by participants to make changes to the access-control policy. All the albums are listed along the left; user groups are listed along the top of the grid; and view, edit, and add permissions are shown as icons in the central grid. This interface also contains a legend at the bottom left.

51

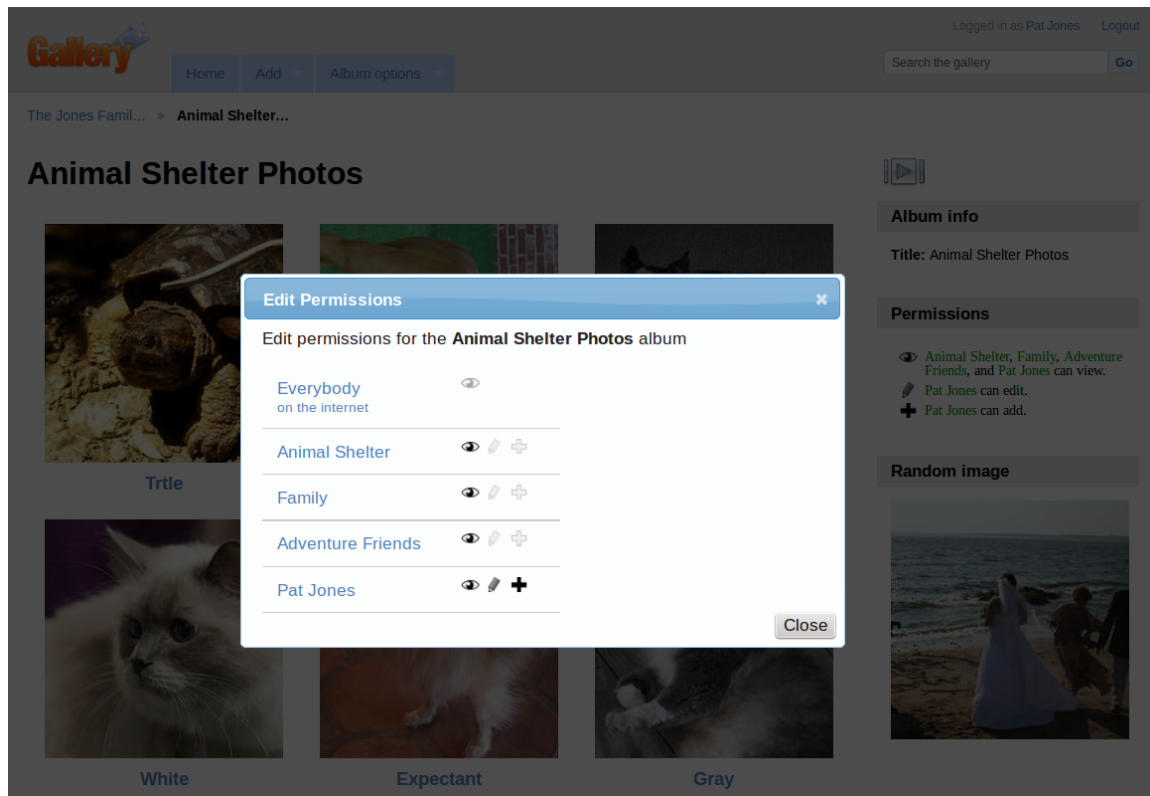


Figure 5.4: Permission-modification dialog. Sentence at the top of the dialog reminds users what album the permissions refer to. The group names are listed along the left side, followed by the different actions (view, edit, and add) that are allowed or denied for that group. A black icon indicates that the permission is allowed; a light grey icon indicates that the permission is denied. Placing the mouse over any icon produces a tool tip indicating the meaning of the current icon. For example: “Animal Shelter can view this album.” Clicking on an icon toggles it between allow and deny.

5.4.3 Conflict resolution and effective permissions

By default, Gallery 3 shows users the implemented-policy rules, but this makes it very difficult for users to accurately understand why a particular group has access, or how to change the permission. When we designed the proximity displays and permission-modification interfaces, we decided to show the user effective permissions (the result of evaluating all relevant implemented-policy rules) rather than the sets of policy rules that induce them. Prior work by Bauer et al. [12] and Maxion and Reeder [66] has shown that people better comprehend access-control policies when they are shown *effective permissions* than when they are shown a list of the implemented-policy rules. In our design we show users effective permissions, and allow them to change permissions by indicating the effective permission they wish to change.

Albums in Gallery 3 can contain subalbums, and permissions on the parent album affect its children, but can be overridden by the permissions set on the children. Similarly, Gallery 3 has two built-in groups: Everybody, and Registered Users. The group Everybody includes all users with accounts on the website and all guest users. The group Registered Users includes all users who have an account on the website. We decided that the Registered Users group caused needless confusion and removed it for the evaluation studies. The existence of these groups effectively makes groups hierarchical, because permissions set on these built-in groups affect the permissions on the other groups.

To address potential permission rule conflicts, Gallery 3 uses the conflict resolution strategy. A conflict can occur any time two or more rules apply to the same user group, album, and action, but have different outcomes (allow, deny). When this type of conflict happens, the system must decide which outcome to use. In Gallery, if users cannot view a parent album, then they cannot view any child albums regardless of the permissions on the child; however, if they can view the parent, then they may or may not be able to view the child based on the permissions on the child. Conflicts in the group dimension always result in a decision of allow. If a user is a member of any group which can view the album, then the user can view the album regardless of the rules on other user groups to which she may belong.

Displaying effective permissions on the proximity displays and permission-modification interface worked well: participants understood the current permission state. Enabling intuitive permission-modification was more challenging because participants were attempting to manipulate effective permissions, rather than specifying rules in which access control is implemented in Gallery. We addressed this issue by translating users' effective permission-change requests into sets of rule changes.

When participants indicated that they would like to change an effective permission, toggling it from deny to allow or vice versa, our algorithm computed the set of rule changes necessary to produce the least number of effective permission changes. For example, assume the album Animals, which has subalbums Dogs and Cats, was not visible to the group Family. The user indicates that she would like Cats to be visible to the group Family. Our algorithm would add an allow rule for (Cats, Family), and (Animals, Family) in order to give Family the ability to view Cats (the parent album must also be visible to Family). The algorithm would also create a deny rule for (Dogs, Family) to ensure that the effective

permissions on the subalbum Dogs do not change.

Participants seemed to find these side effects intuitive. Some participants were briefly surprised when they clicked an icon on the full-page interface and more than one icon changed. However, most participants quickly realized why the change had occurred and did not seem bothered by it. Participants who saw the dialog interface only noticed that multiple permissions were changing when they tried to remove permissions from a group, and the permissions were also removed from the group Everybody. However, similar to the full interface, they rapidly determined the reason for the change.

Chapter 6

Detailed methodologies

To evaluate the proximity display plug-in, we conducted a pre-study and three studies: an eye-tracker study, a lab study, and an online study. All of these studies were designed to test the same three core hypotheses (listed below). However, each study looked at these hypotheses using a different methodological approach and collected different types of data. Individually these studies tell us how people interact with proximity displays, but together they help form a cohesive understanding of user interactions.

In Chapter 7, we present the results of the eye-tracker, lab, and online studies together to give the reader a more holistic understanding of how users interact with proximity displays, and the effect showing these displays has on user behavior. The methodologies for these studies are presented together in this chapter. The pre-study served only to inform the methodological design of later studies, so we mention it in this chapter only to motivate methodological decisions and do not mention it at all in Chapter 7. In this chapter we present the methodologies for the eye tracker, lab, and online studies.

In all of our studies, participants were asked to role play the part of Pat Jones, who manages several online photo albums using Gallery 3, modified to include one of our proximity displays or control user interface. Participants were given some background information about Pat and a series of messages asking Pat to perform various photo management tasks that did not explicitly involve access control. However, participants were told that Pat was responsible for fixing all errors in the photo albums and informed of the relevant access control policies. We designed the photo albums used in this task so that they would contain a variety of errors, including spelling errors in the photo captions and access control errors. We observed whether participants noticed and corrected these errors.

The eye-tracker study was a between-subjects lab study in which an eye tracker was used to better understand when participants were looking at the proximity displays. Tasks were given to participants in a fixed order. During this study we noticed several unanticipated user behaviors (further discussed in Chapter 8) that we then explicitly tested for in the following study. Excluding results from the eye tracker, all results from this study are replicated with greater precision and power in later studies. Consequently we focus on the analysis of the eye tracker data when discussing this study.

The lab study was a between-subjects study conducted in a lab environment. Tasks in this study were presented in a random order and were randomly paired with permis-

sion errors to ensure that permission checking behavior was being measured separately from the influences of task wording or ordering. This study focused on collecting detailed observations of participants and post-study interviews.

The online study was conducted on Amazon’s Mechanical Turk. This study was a within-subjects study where each participant saw both a control condition and an experimental condition. There was a fixed task order and participants had a set time to work on each task. This study focused on evaluating the effectiveness of proximity displays with a large number of participants.

In the following sections, we detail the methodologies and data analysis for all three studies.

6.1 Hypotheses

All the evaluation studies were designed to primarily test the following three hypotheses in a photo sharing system where participants were treating security as a secondary task:

- H1: Correcting/checking permissions** Users who see permission information on a proximity display check and correct errors more often than users who see permission information on a secondary page.
- H2: Permission recall** Participants who see permission information on proximity displays can recall those permissions better than participants who see permission information only if they click to a second page.
- H3: Negative effects** Participants who see proximity displays take no more time, and correct no fewer non-permission errors, than participants who see permission information on a secondary page.

6.2 Eye-tracker study

The eye-tracker study was a 1.5-hour laboratory study in which 34 participants were divided into three conditions: two proximity display conditions and a control condition. In the study, users took part in a role-playing scenario in which they performed a variety of tasks, including various permissions-management tasks on a set of albums. We arrived at the final design for the study after a 4-person pilot.

6.2.1 Study conditions

We tested three conditions: a control and two locations where the proximity display could be located. The control condition had no implemented policy information visible on the main interface, but did have a link to the full-page policy modification interface. The sidebar condition (Figure 6.1(b)) included a proximity display in the sidebar, and the under-photo condition (Figure 6.1(a)) included a display that appeared under each photo or album thumbnail when the mouse cursor was over the photo/album thumbnail.

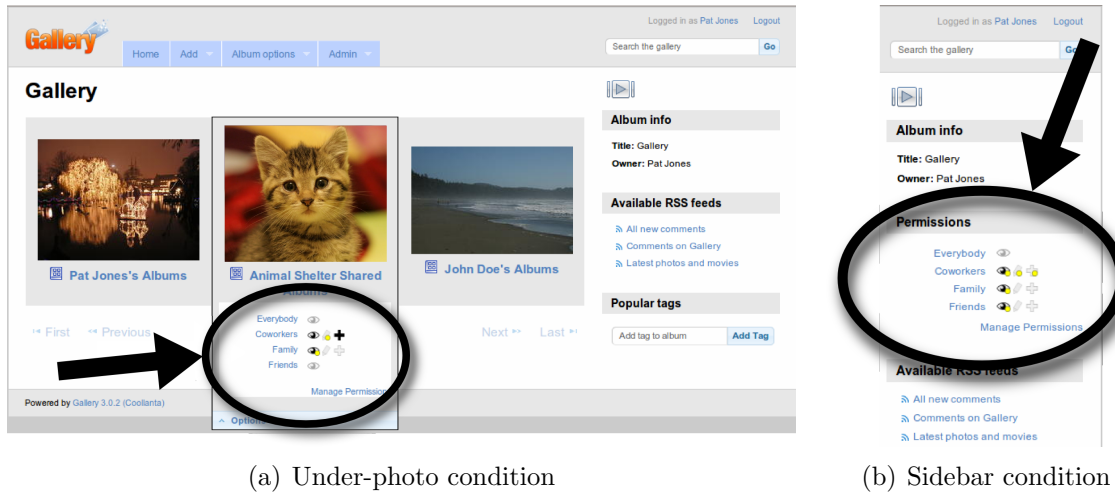


Figure 6.1: The two proximity display conditions used in the eye-tracker study: (a) in the under-photo condition, and (b) the sidebar condition. Proximity display in the under-photo condition. The proximity display in (a) shows that the group Everybody has no permission; Coworkers can view and add to this album and all subalbums, but can only edit some subalbums; Family can view this album and some subalbums; and Friends cannot view anything.

6.2.2 Protocol

The eye-tracker study was a between-participants design with a round-robin assignment to experimental conditions. A think-aloud protocol was used. Participants in all conditions performed the same tasks, and the only variable between conditions was the Gallery 3 interface participants were exposed to. The tutorial used to familiarize participants with the Gallery 3 interface also differed slightly by condition.

Participants were asked to role play the part of Pat Jones, who manages several online photo albums using Gallery 3. During the course of the study, participants received information about events in Pat's life, including emails from coworkers, family, and friends. These emails, delivered to participants in printed-out form by the researcher administering the study, included requests from Pat's coworkers, family members, and friends to perform various tasks with the online albums.

As Pat Jones, participants started with a tutorial that asked them to walk through manipulating photos using Gallery 3 that had been previously set up with seven albums in hierarchies and simplistic permissions. When participants completed the tutorial, the researcher had them open a new Gallery 3 site that had many more albums and more complex permissions. These albums did not overlap the tutorial albums.

After the tutorial, participants were first asked to perform five clearly defined and progressively more complex warm-up tasks (rows 1–5 in Table 6.1): rotate a photo, read a permission, delete a photo, change a permission, and change some titles. If any tasks were not successfully completed, the researcher prompted participants with an email that pointed out the error; if participants still could not complete the task, they were verbally

Task	Area	Permission error	Album state	Prompted
Work Information Page				
1-5	Warm-up	Read, Add	Existing	Prompt
6	Coworkers	None	Existing	None
7	Coworkers	Add	New	None
8	Coworkers	Remove	Existing	None
9	Coworkers	Read	Changed	Prompt
Friends Information Page				
10	Friends	Remove	New	Prompt
11	Friends	Read	Existing	Prompt
12	Friends	None	Existing	None
13	Friends	Add	Changed	Prompt
Family Information Page				
14	Family	Add	Existing	Prompt
15	Family	None	Existing	None
16	Family	Read	New	None
17	Family	Remove	Changed	Prompt

Table 6.1: Tasks and information given to eye-tracker study participants.

instructed by the researcher how to do so. This was done to ensure that all participants knew how to operate Gallery 3 and to help them get acclimated to working with the albums.

The bulk of the study consisted of tasks 6–17, summarized in Table 6.1. Each task was composed of a set of *actions*: individual permission, rotation, deletion, spelling, or re-naming errors that needed to be corrected. Each task had a primary action directly expressed by the email sender and several additional actions implied by errors such as rotated photos. Some tasks conveyed the ideal policy to participants using text in the email (shown in the third column of Table 6.1). The tasks were divided into three sets based on whether the albums participants would manipulate contained photos of coworkers, friends, or family (shown in the second column of Table 6.1). Before each set of tasks, participants were given an information sheet explaining Pat’s normal interactions with this group of people. Half the tasks required adding or removing a permission (shown in the third column of Table 6.1). A quarter conveyed to participants desired permissions, but no permissions needed to be changed. The final quarter had no access-control component. All tasks contained at least one title, rotate, delete, or organize action intended to distract participants. Each task was performed on albums in one of three states (shown in the fourth column of Table 6.1). *Existing* albums were already set up in Gallery 3 when participants started. *New* albums were created by participants. *Changed* albums were those that participants had previously read or changed a permission, but, unknown to participants, some part of the album had been altered by the researcher after participants had last seen the permissions. Tasks for which failure to correct a permission error resulted

in an email calling this out were called *prompted*; all others were *unprompted* (rightmost column, Table 6.1). When participants failed to complete a prompted task they received an email from one of Pat’s coworkers, friends, or family members pointing out the error and requesting that it be fixed.

In addition to the task-related albums, there were four albums that participants were never directed to interact with. Two of these albums had correct permissions and two albums had incorrect permissions.

At the end of the study, participants filled out a survey that asked them to recall the view and add permissions for every album they worked with, the two albums that had incorrect permissions but were not part of a task, and two non-task albums with correct permissions. For each suggested combination of album, group, and permission participants could answer *True*, *False*, or *Not Sure*. For each set of questions about an album, participants were asked how confident they were of their answers.

6.2.3 Recruitment and demographics

We recruited 34 participants using a university-run electronic bulletin board for advertising research studies. Participants ranged in age from 18 to 41 with a mean age of 23.9. Twenty two participants were students. One participant was excluded due to an inability to complete even half the study in the allotted 1.5 hours. After this exclusion, we were left with 11 participants per condition.

6.2.4 Data collection and analysis

We collected and coded data derived from a combination of in-session notes, screen-capture video, audio, exported information from an eye tracker, a snapshot of the resulting permission state of the photo website, and the survey. All data was loaded into a database so information from different sources could be correlated.

Eye tracker

We used an SMI eye tracker to record video of events occurring on the screen, audio of participants, and the time and screen coordinates of fixations and user events (e.g., mouse clicks).

In the under photo condition, proximity displays appeared below photos and tended to be visible for only short times. To determine when and where displays appeared on the screen for each user, we used a custom Matlab script that scanned each video frame for a unique static part of the proximity display and recorded the time and location of each display. This information was then matched with the fixation data from the eye tracker to determine when and where participants saw proximity displays.

6.3 Lab study

The lab study was a 1.5 hour between-subjects study where 33 participants were divided into four study groups based on two treatment types: proximity display and permission-modification design. Half of the participants saw permission information on a proximity display located under every photo and album thumbnail, and the other half saw no proximity display. Similarly, half of the participants modified permissions using a full-page permission-modification interface (Figure 6.3(a)), and the other half modified permissions using a pop-up dialog (Figure 6.3(b)). We arrived at the final design for the study after a 17-person pilot.

In the eye-tracker study we observed that participants were more likely to check permissions on some tasks than others. Based on think-aloud data we hypothesized that the wording of the task was impacting the permission checking behavior. To address this, in the lab study we made the wording more consistent across task emails. We also randomized which tasks had permission errors associated with them. Similarly, conveying the ideal policy in the email text in the eye-tracker study may have influenced participants' permission checking behavior. To address this we presented the ideal policy to participants all at once, instead of in each email.

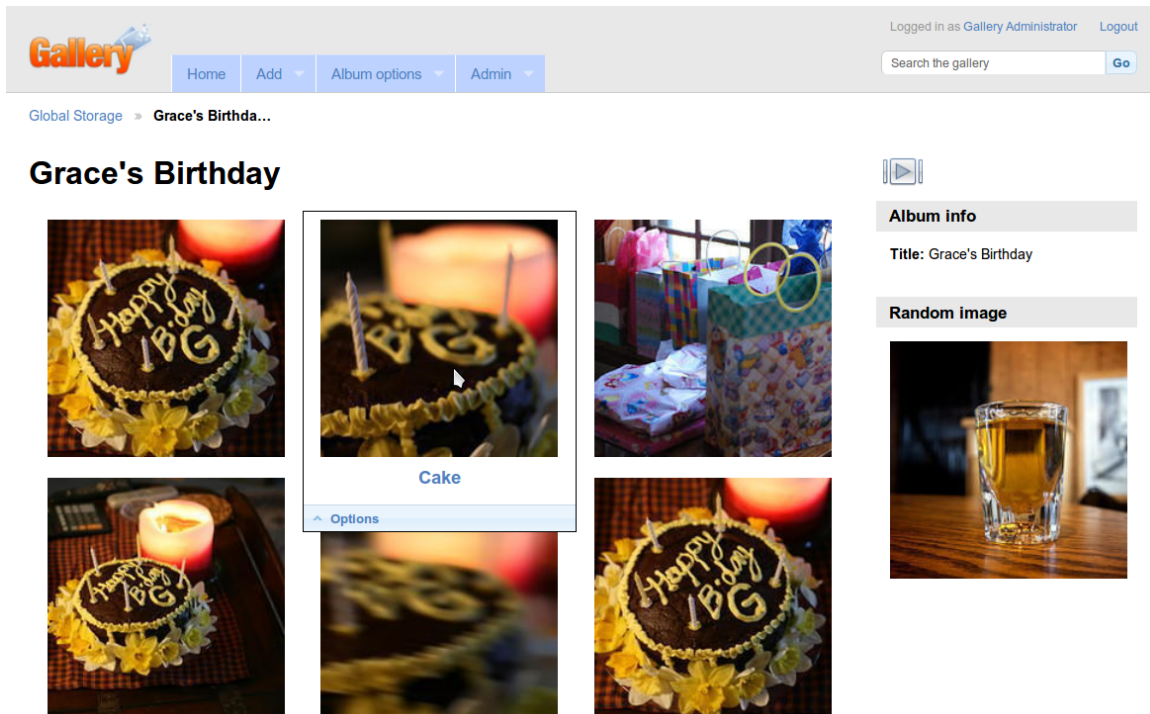
6.3.1 Study conditions

This study had two experimental variables: proximity display, and permission-modification interface. Both variables had two levels, resulting in four experimental conditions:

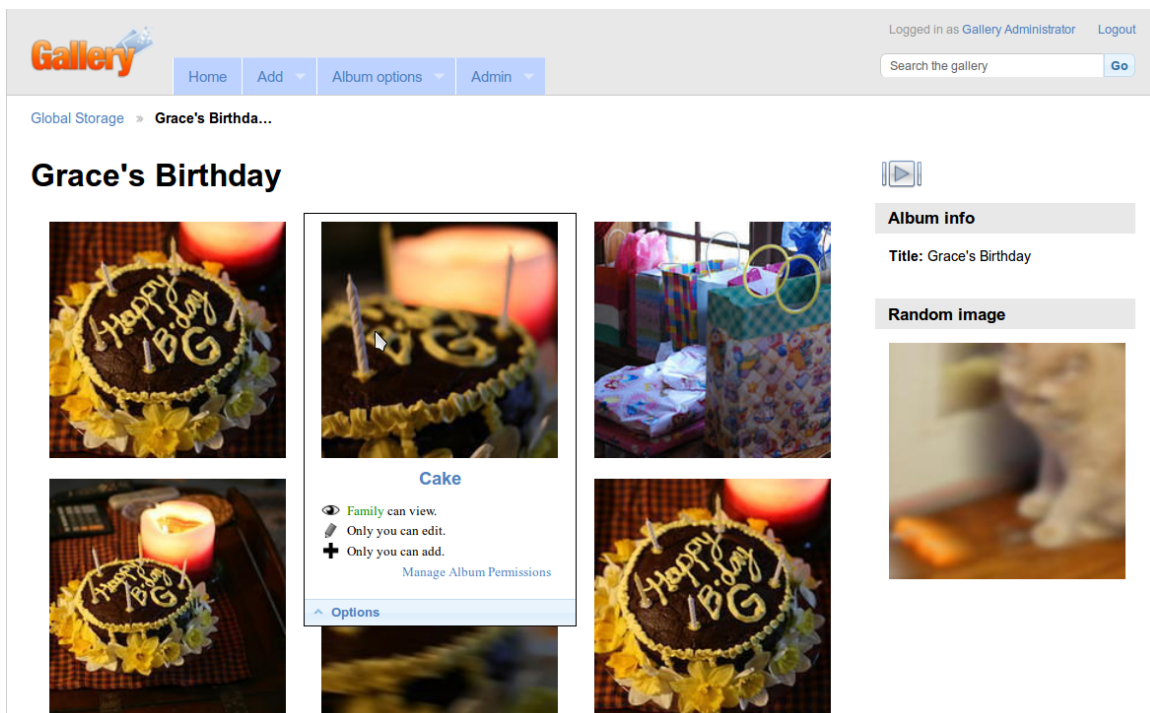
Condition name	Proximity display	Permission modification interface
Control Dialog	None	Dialog
Control Full	None	Full page
Under Dialog	Under photo/album thumbnail	Dialog
Under Full	Under photo/album thumbnail	Full page

Proximity display – Participants in the control condition see no permission information on the photo management interface (Figure 6.2(a)). To access the permission-modification interface, control participants must select “edit permissions” from one of the options menus. Participants in the under-photo condition had the option of placing their mouse over an album or photo thumbnail to see the proximity display (Figure 6.2(b)), or using the “edit permissions” link in one of the options menus.

Permission-modification interface – When participants clicks on any of the “edit permissions” links or uses the “manage permissions” link on a proximity display, they are taken to a permission-modification interface. Participants in the dialog condition see a permission-modification dialog that allows them to view/modify permissions for this album only (Figure 6.3(b)). Permission information for other albums is not shown on the dialog. Participants in the full permission-modification interface condition are taken to a new page where they can view/modify permissions for any album (Figure 6.3(b)). To assist users, the album they were previously viewing is highlighted.

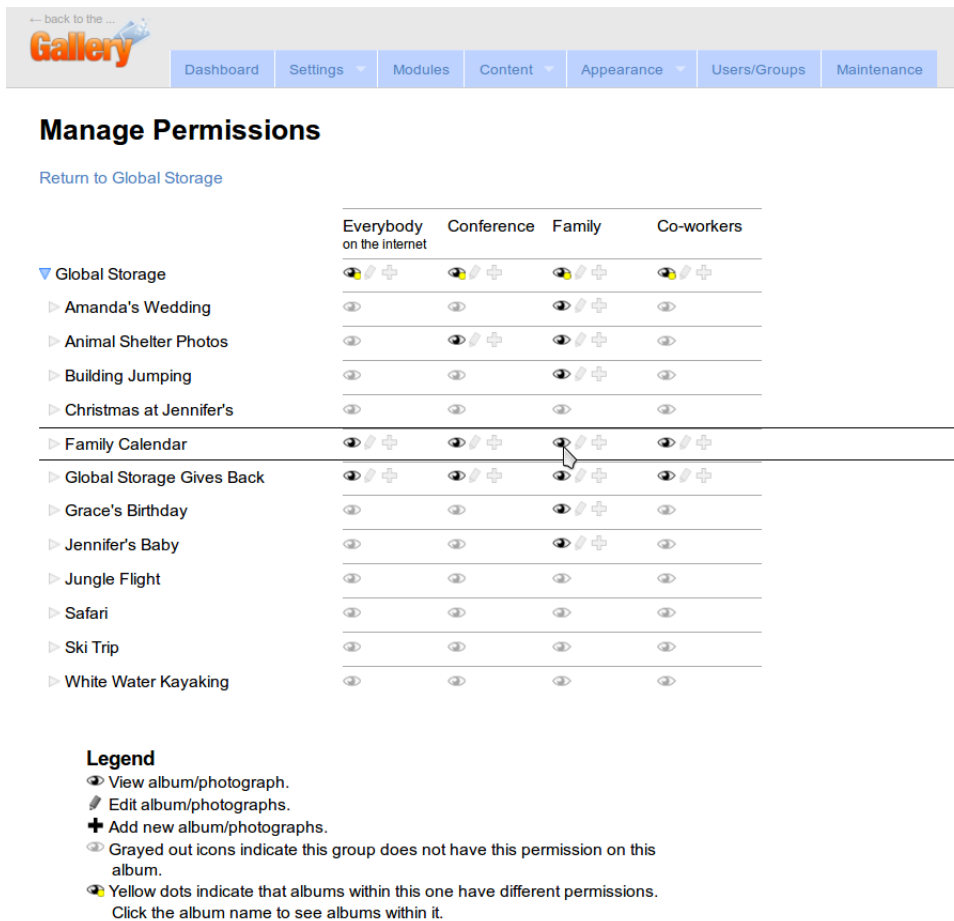


(a) Control

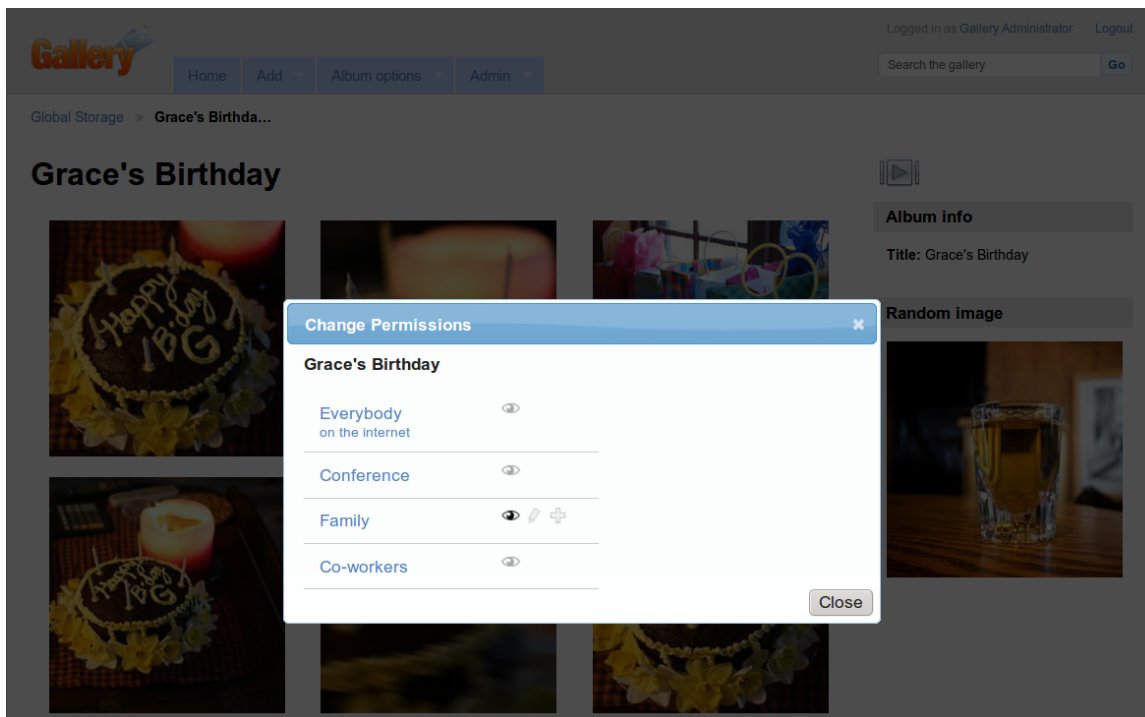


(b) Under photo

Figure 6.2: Gallery 3 interface without a proximity display (a), and with a proximity display under every photo and album (b).



(a) Full-page permission-modification interface



(b) Dialog permission-modification interface

Figure 6.3: Full-page permission-modification interface (a) and dialog permission-modification interface (b).

1. Tutorial
2. Instructions about Pat and ideal policy
3. Training
 - (a) Task 1 (Rotate, delete, cover, spelling, change permissions)
 - (b) Task 2 (Move, read permissions)
4. Prompts – feedback on both training tasks
5. Three policy comprehension questions
6. Three tasks randomly drawn and removed from the set of tasks with permission error drawn and removed from the set of permission errors and non-errors such that 50% of tasks have errors
7. Task 4 with permission error randomly drawn and removed from the set of permission errors
8. Task 4 task prompt
9. Eleven tasks randomly drawn and removed from the set of permission errors and non-errors such that 50% of tasks have errors
10. Memory and comprehension questions
11. Debriefing interview

Figure 6.4: Lab study protocol order.

6.3.2 Protocol

The lab study was a between-subjects design with a round-robin assignment to experimental conditions. A think-aloud protocol was used. Participants in all conditions performed the same tasks and saw the same permission errors. However, tasks and errors were shown in a random order.

Participants were first asked to read and fill out a consent form. Next participants were given the opportunity to interact with the eye tracker. The pre-study indicated that participants who understood where the eye tracker could and could not see them were more likely to stay in range during the study. After interacting with the eye tracker, participants were trained in how to think-aloud followed by a short calibration of the eye tracker.

Participants were then presented with a training version of the Gallery 3 site. The training version was identical to the website used in the primary section of the study with the exception that it had a different and smaller set of albums and photos. Participants were verbally given a user name and password and asked to log in. They were then given a printed tutorial and asked to work through it, and make changes to the website as they went. The tutorial clearly stated that this was a practice version of the website and made it clear that experimenting would not impact the main study. The tutorial covers how to navigate Gallery 3, move photos between albums, change titles on photos and change permissions. Based on the pre-study, we decided not to train participants on how to rotate a photo or change the cover image on an album as both of these features are easily found on the same menu as title manipulation, which the participants are already trained to find. Participants took an average of 5 minutes 27 seconds to complete the tutorial.

Participants were asked to role play the part of Pat Jones, who manages several online photo albums using Gallery 3. During the course of the study, participants received emails from coworkers. These emails, delivered to participants on paper by the researcher administering the study, were requests from Pat’s co-workers to perform various tasks with the online albums. Participants were allowed to look back through any piece of paper given to them, including the tutorial and the instruction sheets. This was done so that participants could perform study tasks without having to memorize all of the instructions.

Gerald's Photograph Policy

1. No photographs containing drugs, alcohol, or anything inappropriate.
2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.
3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great a company we are.
4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.
5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

Figure 6.5: Gerald's Photograph Policy

Participants next opened the main Gallery 3 website, and were given two instructional papers. The first paper described Pat Jones, an employee of Global Storage responsible for maintaining the online photo album. The second instructional paper explained that part of Pat's job involved helping other co-workers with the photo system. However, it was also Pat's job to enforce the boss, Gerald's, photo policy (Figure 6.5).

The bulk of the study was comprised of 2 warm-up tasks and 14 normal tasks. Tasks in this study are composed of a two paragraph email and an associated album. Each email started with a short paragraph from the sender describing the album, that clarified whether the album was personal or professional. The second paragraph named the album participants were to work with and stated the set of *explicit errors* the sender would like Pat to complete. The album also contained at least one photo that conflicted with Gerald's rules; we refer to these conflicts as *implicit errors*. Each album contained either personal or professional photos and 50% of the tasks were associated with personal albums. Task and permission error orders were randomized so as to remove effects caused by task wording or ordering. The only exception was the fourth non-warmup task, that always occurred in the same location and always had a randomly selected permission error. This task was a prompted task designed to make users feel like someone was checking their work. It had the same task and always had an error for consistency between participants. In this way 50% of the personal tasks had permission errors and 50% of the professional tasks had permission errors, but the errors occur in a random order and were paired with a random task. Tasks began when participants were given the associated printed email and ended when participants requested the next email.

Participants were initially given two warm-up tasks, one-at-a-time, that matched the description of tasks in the prior paragraph in every way except that they had fixed permission errors that were never randomized. The two training tasks explicitly instructed participants to perform every action needed in the study:

- Rotate a photo
- Delete a photo
- Change an album cover

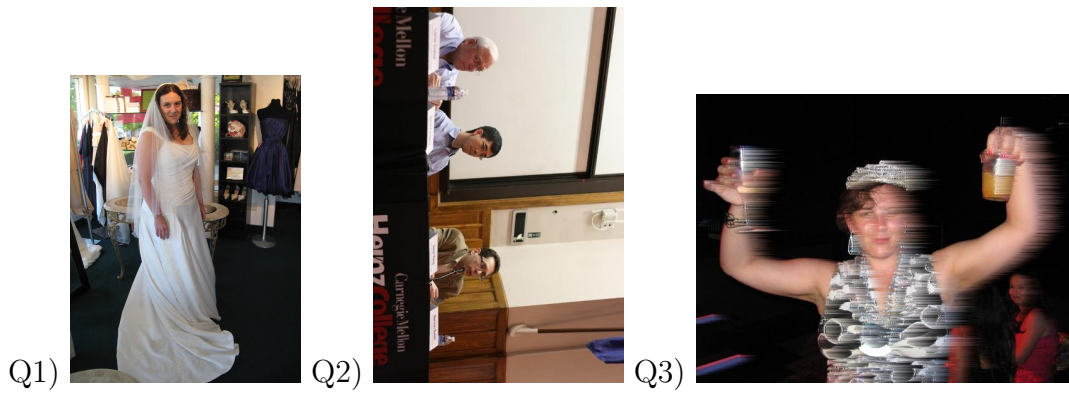


Figure 6.6: Participants were asked, by a co-worker, whether each of the above images were acceptable to post on Gallery 3 and whether the co-worker needed to make sure to do anything when they put the photo on Gallery 3. Q1 has no problems with the photo, but should be visible to Friends and Co-workers only. Q2 needs to be rotated and should be visible to Everybody on the Internet. Q3 cannot be uploaded to Gallery 3 because it is blurry and contains alcohol.

- Move a photo
- Change a title
- Change a permission

Permissions were never expressly mentioned in any of the emails, so for consistency they were not mentioned in the warm-up emails. To be consistent with the permission errors, the other types of errors mentioned in Gerald’s photo rules were never expressly mentioned in any email. Instead of mentioning them in the warm-up, we asked participants to delete a photo for another reason. If participants had technical trouble completing an action during the warm-up then the researcher provided additional instruction. If participants did not correct an error, then, after both training tasks, they were sent a prompting email from their boss. This email clearly stated that he noticed that participants had not performed the action correctly and asked participants to fix the issue. This prompt clearly stated what error participants had failed to correct.

Participants next received three emails, one-at-a-time, from a co-worker with three different photos attached (Figure 6.6). The co-worker asked whether the attached photo was acceptable to post on Gallery 3 and if so, was there anything the co-worker needed to make sure to do. These emails were used to be certain that participants understood and could apply Gerald’s policy. If participants did not mention permissions when responding to the emails they were sent an email from the co-worker asking whether they needed to do anything with the permissions. If participants answered the question incorrectly, they were sent an email containing Gerald’s rules from their co-worker.

Participants were next given three tasks, as described above. These tasks were randomly ordered and were *unprompted* – If participants failed to correct any of the explicit or implicit errors, no action was taken and they were allowed to move on to the next task with no prompting or reminders.

The fourth task was not randomized and every participant saw the same task at the same point in the study. The fourth task contained two rotation errors, two alcohol errors, and a permission error, which was randomly selected from the set of all personal permission errors. If participants failed to correct any of the explicit or implicit errors, they received a prompting email from Gerald, their boss. The email said that Gerald checked the album and was disappointed in Pat. The prompting email also contained a list of Gerald’s rules, but did not specify what Gerald found to be incorrect. Based on the pre-study we found that the fourth task was the point where participants started realizing that no one was checking their work. It was also the point where they were experienced with the interface but may have forgotten about the rules. By providing a prompt here, we ensured that participants remembered to enforce the rules and realized someone was going to be checking their work. If participants failed to correct all the errors after the prompt, no action was taken and they were allowed to move on to the next task.

The remaining ten tasks were presented to participants in random order. If participants failed to fix any errors in these tasks, no action was taken and participants were allowed to move on to the next task without any intervention.

After completing all tasks, the researcher asked participants a set of recall questions for the last four albums and the third and fourth albums they saw. Participants were asked to recall what the permissions were on those albums the last time they interacted with them, as well as what permissions the album should have had according to Gerald’s policy. Participants were then asked to recall, in their own words, what Gerald’s policy was.

Finally, the researcher engaged participants in an unstructured interview to better understand the decisions and behaviors they had engaged in. Many of the behaviors were short (1-2 seconds long), so we used a contextual interview approach [40] to help people remember what they were doing when they engaged in the behavior. For each question, participants opened the album they had been working with and the researcher explained the context that the behavior occurred in and asked participants questions concerning what they were thinking or what they had been trying to do. The researcher also asked about prior experiences and opinions that might have impacted participant behavior.

6.3.3 Participants

Participants were all native English speakers who had previously uploaded photos to an online social network or photo-sharing website. They were recruited using an existing pool of people interested in participating in behavioral research studies. This pool includes both students and members of the Pittsburgh community. They ranged in age from 18 to 53, with a median age of 22. Participants were predominately students (60%) and female (66%). All participants had previously shared photos online using Facebook, and the majority had also shared photos using other websites. No lab study participant dropped out of the study, but one participant was excluded for not completing all the tasks in the time allotted.

6.3.4 Data collection and analysis

Participants were audio recorded and a screen capture program recorded a video of their web browser. Our custom version of Gallery 3 recorded detailed logs of the participants' actions in a database. The researcher took detailed notes during the session, including a timestamp whenever she handed participants a paper. Data collected from participants falls into seven categories: permission correcting, permission-checking behavior, action order, non-permission correcting, permission recall, rule comprehension, and interview data.

Log data

We altered Gallery 3 to log the following activities, every log entry contained a timestamp and the Gallery 3 assigned ID of the affected album or photo:

Webpage navigation: We logged the URL of every page participants navigated to.

Photo-manipulation actions: Whenever participants performed a photo- or album-manipulation action we logged the timestamp, the id of the photo or album, and the value that was changed. Photo-manipulation actions involved changing an album cover, deleting a photo, moving a photo, changing the permissions, changing the title, or rotating a photo.

Proximity-display contents: On every page load we recorded the HTML and spatial position of every proximity display on the screen.

Proximity-display visibility: Displays were only visible if users placed their mouse over an album or photo thumbnail. We logged every time a mouse was placed over a display and every time the mouse moved off of the display.

Permission correcting

In the Gallery 3 system, permissions are associated only with groups and albums. Neither photos nor users can have permissions associated with them. Permissions are described as a trio of user, album, and action. Each user, album pair has three actions associated with it: view, add, and edit. In this study we tested seven different permission errors, that are shown in Table 6.2.

At the end of every study session, we archived the state of the final permission settings and automatically extracted them into a database. Each permission was compared to its initial state and marked as “changed” or “unchanged.” The permission was also compared to its correct state and marked as “correct” or “wrong.”

Permission-checking behavior

The researcher recorded in their notes every time participants *explicitly checked* a permission. To approximate when participants actually notice permission errors, we measured *explicit checking* behavior. Control participants are said to have *explicitly checked* permissions if they opened the permission-modification interface. Participants who were shown proximity displays are said to have *explicitly checked* permissions if they (1) opened

Professional/Personal	View	Add	Edit
Professional	Nobody can view	-	-
Professional	Family and Co-workers	Co-workers	Co-workers
Personal	Family, Co-workers, and Conference	Co-workers	Co-workers
Personal	Family only	Nobody	Nobody
Professional	Family and Co-workers	Co-workers and Family	Co-workers
Personal	Everybody on the Internet	Co-workers	Co-workers
Professional	Family and Conference	Nobody	Nobody

Table 6.2: Implemented policy errors. Each participant experienced every error once during the 14 tasks. The first training task had a permissions error where everybody on the Internet could see a personal album, so participants would have seen this error twice during the study session.

the permission-management interface; or (2) read the permission aloud; or (3) indicated through mouse behavior that they were reading the permission display (moving the mouse under the words while reading or circling the display with the mouse); or (4) pointed at the permission display with their hand while staring intently at the screen. The identification of explicit permission-checking behavior in the under-photo condition was done by the researcher during the study and recorded in the researcher’s notes, that were later coded in the database.

The majority of participants explicitly checked permissions during the task involving the album associated with that permission. However, some participants checked all permissions at once, either by looking at permissions one-by-one on the full permission-modification interface or by mousing over each album thumbnail one-at-a-time, reading the proximity display and correcting the permission when wrong. These participants were flagged in the database as having *checked permissions all-at-once*. Unless otherwise noted, the analyses in the results section treat participants who are flagged as having checked permissions all-at-once as having checked permissions on every task.

The researcher’s coding for checking permissions was compared with the action logs to determine approximate accuracy. For the control participants, the action logs and the notes were perfect matches. Control participants are forced to open a permission-modification interface to check permissions, and therefore all permission checking behavior is recorded in the action logs. Under-photo participant action logs were also compared to researcher notes; the action logs were a subset of the researcher notes indicating that the researcher had not missed any programmatically measurable permission-checking behavior.

Action order

The Gallery 3 software was modified to record detailed logs of all participant actions. These logs include the type of action performed, what photo/album it was performed on, any change in value, and a timestamp. After the study, the action log was annotated with the task participants were engaged in at the time they performed each action. The task information was determined based on the timestamps the researcher manually recorded at the beginning of each task.

For each user and task, we analyzed the logs and created a list of the minimum and maximum timestamps for each action type (change cover, delete, move, permissions, rename, and rotate). If users never engaged in a particular action type during a task, i.e. no timestamp existed, then the action was excluded from the list. We then ordered the actions based on the minimum timestamps and coded each action as *first*, *middle*, *last*, or *only*. We refer to this ordering as when the action was first engaged in. Similarly, we ordered the actions based on the maximum timestamps and coded each action as *first*, *middle*, *last*, or *only*. We refer to this ordering as when the action was last engaged in. If participants only engaged in one action on that task, the action is marked *only* as there can be no order with only one action.

For example, suppose that a participant performed three action events during a task: rotate, delete, and rotate. Then, the first engaged in ordering is rotate and delete, where rotate is first and delete is last. Conversely the last engaged in ordering is delete then rotate.

We chose the *first*, *middle*, *last*, and *only* codes based on the observation that in 70% of the tasks participants engaged in three or less types of actions with an average of 2.8 different action types. Participants were free to take as long as necessary to complete tasks, leading to a high variation amongst participants. We chose to code the action order rather than normalize the timestamps because we felt it provided a more accurate picture of the order that participants engaged in actions.

Non-permission error correction

Gallery 3 stores the meta-data of all photos and albums in a MySQL database. After each study session, several scripts were run against this database to collect and code relevant information and then archive the database. The scripts compared the meta-data on all albums and photos to the default meta-data. If the default value was different from the final value, then the album or photo was marked as having been rotated, re-titled, deleted, cover changed, or permissions changed depending on the meta-data. A second script was run that coded each meta-data element as *error*, *error fixed*, or *no default error*. For example, some tasks required participants to make any change to the title of a photo. For these tasks, any change to the title was considered correct. Some tasks required participants to make a specific change; in these cases the applicable meta-data value had to match the required final state.

Permission recall

At the end of the study, participants were verbally asked by the researcher to recall the final permissions of the last four albums they saw (tasks 10-14) and the albums from tasks 3 and 4. The last four albums were chosen because they were the most recent and therefore the most likely to be recalled. The album from task 3 was chosen because it was seen less recently, and before the prompt on task 4. The task 4 album was chosen because participants were prompted to change permissions on that task and were more likely to have seen and interacted with them. The participants' answers were coded such that they

could be directly compared with the actual permissions in the database.

Rule comprehension

After participants completed both training tasks they were asked three comprehension questions. The questions were given to participants one at a time via an email from a co-worker. The emails asked participants whether each of three photos matched Gerald’s policy and if not, what changes needed to be made. Participants’ answers to these questions were coded in terms of 1) whether the photo was suitable to be placed on Gallery 3 at all, and 2) which, if any, of the existing policy violations participants mentioned.

At the end of the study, in addition to being asked to recall permissions, participants were asked to apply Gerald’s policy about permissions to each of six albums. Participants were asked what permissions Gerald would have wanted each album to have. These answers were coded and compared to Gerald’s policy

Unstructured post-interview analysis

The post interviews with participants were transcribed and question and answer pairs were broken into topics. Each topic was printed on a slip of paper that was used in an affinity diagram [17].

6.4 Online study

The online study was an hour long within-subject study with 658 participants and 5 treatments. Each treatment was composed of a control condition and an experimental condition, that were shown to participants in serial. Participants took part in two different role-playing scenarios in which they performed a variety of tasks, including various permissions-management tasks on a set of albums.

The goal of conducting an online study was to test the proximity-display interface on more participants than could feasibly be brought into the lab. By using Mechanical Turk, we were able to get a large number of users in a relatively short time frame.

In the prior studies we noticed that some participants were more inclined to check permissions than other participants. This made us concerned that the inter-subject variability was high and might be impacting our results. To account for this potential issue, we made the online study a within-subjects study, so we could measure how the same participant performed with and without proximity displays.

6.4.1 Study conditions

Participants in the online study were assigned round robin to one of five treatments, two condition orders, and two scenario orders, effectively assigning them to one of 20 possible treatment combinations. Each treatment was composed of two conditions: a control condition showing no access-control information, and an experimental condition displaying a version of the proximity display. It was also composed of two scenarios: a work scenario

where Pat has to manage work related albums, and a home scenario where Pat has to manage family and friend albums.

To prevent biasing either condition we ensured, via a round-robin assignment to all 20 treatment combinations, that participants were equally likely to first encounter the control condition as the experimental condition. We also ensured that the home scenario was equally likely to appear first as the work scenario, and that they were equally likely to be paired with the different conditions.

There were five experimental treatments in this study:

Audit – The audit condition showed, in the proximity display, who had recently accessed the album and what groups they were in. This display was visible under the album thumbnail, and when the album was opened it was displayed on the sidebar. If a user group had the ability to view an album, then we set up the audit display to show at least one person in that group accessing the album recently. Similarly, user groups who did not have access were never shown as having previously accessed the album. For example, the Animal Shelter album should be visible to Animal Shelter Employees and friends, but in our study it was visible only to friends. During the study, the audit display showed that several members of group Friends had viewed the Animal Shelter album but that no members of group Animal Shelter Employees had viewed the album.


Facebook – The Facebook condition was intended to simulate Facebook’s access-control permission indicators as closely as possible. We decided to use Facebook’s user-interface design because it is both a very popular site for sharing photos and its user interface design is very different from our own. Facebook uses a set of icons to express the privacy policy associated with albums. An album can be publicly visible (🌐), visible only to the owner (👤), visible only to friends (👥), or a custom settings (⚙️). Similar to Facebook’s user interface, we placed the relevant icons under each album thumbnail, and when the album was opened we placed the icons in the upper right hand corner. Mousing over the icon resulted in a pop up listing the groups who had the right to view the album. However, clicking on the icon resulted in our permission-modification dialog rather than Facebook’s drop down menu. Since we are testing whether people can notice errors, rather than the impact of the permission-modification interface design, we felt it was more important that the permission-modification interface be consistent across conditions than for it to be consistent with Facebook.

Mixed – The mixed condition showed the proximity display under the album thumbnail, and when the album was opened the proximity display was shown on the sidebar. This condition was selected based on the outcome of the two prior studies showing that participants check permissions at the beginning and ending of tasks, and that participants use the display under the album thumbnail to determine the presence of an error. This condition was expected to both support this behavior and take up less screen real estate than the under-photo condition.

Sidebar – The sidebar condition shows the proximity display on the sidebar. No permission information was ever shown under the album or photo thumbnails. This condition setup is identical to the ones used in the eye-tracker study.

Under Photo – The under-photo condition shows the proximity display under the

album thumbnail, and when the album is opened it shows the display under every photo. This condition setup is identical to the ones used in the eye tracker and lab studies.

We were concerned that the appearance/disappearance of the proximity display when participants switched conditions would draw unnecessary attention to the display, and bias our results. To counter this issue, we created a similar proximity display that showed tag information instead of permission information. This display was placed on the control interface in the same location as the permission information shown in the experimental interface (Figure 6.7). For the audit, mixed, sidebar, and under conditions the tag display simply appeared in the same location as the permission display. For the Facebook condition a tag icon () was displayed in the same place where the permission information icon was shown (Figure 6.7, Tables 6.3 and 6.4).

6.4.2 Participants

Participants in the online study were recruited using Amazon’s Mechanical Turk. Participants coming from Turk were first shown a bulleted list describing what the study entailed, then they were shown the consent form, and if they agreed to it they were assigned a study group and began the study. Table 6.5 shows the number of participants who completed the study in each condition (column 2), changed at least one permission without being told to do so (column 3), and the number of users who agreed to the consent form but did not complete the study (column 4).

Participants came from a wide range of professions and education levels. The most common profession was Student (26.3% of participants), followed by Unemployed (14.1% of participants). Only 5% of participants reported a technical profession. The most common education level was Some College (39.9% of participants), followed by Bachelors Degree (29.0% of participants). Participants ranged from 18 to 63 years of age with an average of 28 years old and 46.9% were male.

6.4.3 Protocol

This study was a within-subject online study conducted on Mechanical Turk. Participants were asked to read instructions, do a training, and complete eight tasks for each of two conditions. After experiencing both conditions, participants were asked to fill out a survey that asked memory questions about both conditions as well as participants’ demographics.

When participants visited the study from Mechanical Turk’s website, they were shown a page warning them that the study would take a full hour and they must complete at least 25% of the task components to be paid. If they chose to continue, they were shown a consent form explaining that this was a photo management study.

The study’s web interface was divided into two frames, as shown in Figure 6.8. The top frame showed instructions and emails to participants. The bottom of this frame contained a control bar that allowed participants to shrink the frame, obtain instruction on Gallery 3’s features, and move to the next task. The lower interface showed the Gallery 3 website participants were currently working with.

Each participant went through two nearly identical sets of tasks, one set per condition. To ensure that the two sets of tasks did not appear identical, we created two Gallery 3 websites: one focused on Pat’s personal life, and the other focused on Pat’s professional life. We refer to whether the site concerned personal or professional photos as the site’s *scenario*. The sites had different themes, titles, and photo content, but were otherwise identical. Care was taken that the albums in both sites had the same number of photos and the same type and number of permission and non-permission errors. The order that participants encountered the two sites was assigned round robin.

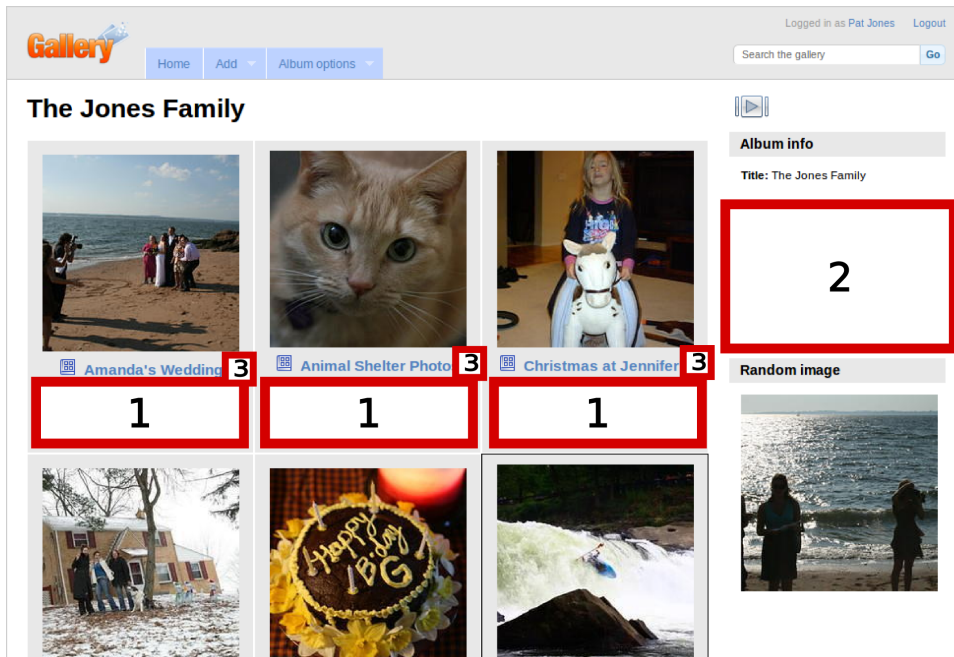
Participants completed six training tasks on each site: open an album, rotate a photo, change a title, change a permission, change a tag, and move a photo. These tasks were expressed as stated directions. For example: “delete the blurry blue teapot.” If participants had trouble completing a training task, they could select one of the instructional pages from the “show me how to” drop down. These instructional pages were viewable at any point in the study, but were most used during training.

After training, participants were shown an instructional page telling them that their personal relations or boss expected Pat to assist in keeping the online photo albums to a certain standard, expressed by the set of rules shown in Figure 6.9. These rules differed only as necessary to accommodate the scenario. Similar to the lab study, this rule set included both permission and non-permission rules. There was one rule about tags, that states which tags need to be present. There were two rules about permissions that state who should and should not have the ability to view albums. The second permission rule is never associated with an error and primarily exists to support the logic of the scenario. Pat should always have access to the albums; we did not want participants thinking they might remove their own access if they removed the group Family.

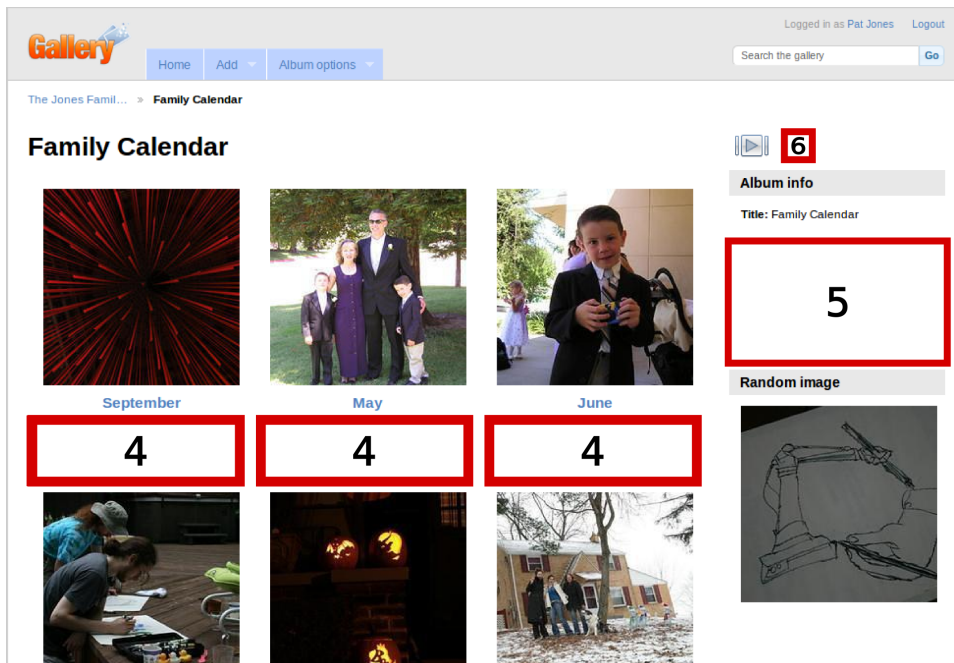
Participants then went through eight tasks one by one. Tasks were introduced by a two-paragraph email shown in the top frame of the browser window. Each email starts with a short paragraph from the sender talking about the album to be worked with; this paragraph makes it clear which group the album is associated with. The second paragraph names the album participants were to work with and contained the set of *explicit errors* the sender would like Pat to correct. The album also contains at least one photo that conflicts with the rules. We refer to these issues with photos as *implicit errors*.

The first task was *prompted*: if participants missed any error, permission or otherwise, they were shown a prompting email pointing out the error and requesting that it be corrected. The remaining seven tasks were *unprompted*; no notification was given to participants if they failed to correct an error. However, we were concerned that some Mechanical Turk participants would attempt to click through the study without making any changes. If participants made no changes during a task, they were prompted by an error message that indicated they had not yet done the task. Participants were given a maximum of 2 minutes to complete each task, each prompting email reset the timer to 2 minutes, effectively giving participants as long as necessary for task 1.

Tasks 3, 5, and 8 were all conducted on the same album. In task 3 there exist both permission and tag errors associated with the album. When participants started task 5, the script automatically adds three photos to the album, and the task email asks participants to interact with these photos. No changes are made to the permissions or tags, effectively



(a) Gallery 3 albums page. Displays title and thumbnails for all albums.



(b) Gallery 3 photos page. Displays title and thumbnails for all photos located inside a single album.

Figure 6.7: Gallery 3 interface showing all the albums and their cover thumbnails (a), and the interface showing all the photos contained within a single album (b). Proximity-display locations are marked with numbers 1-6 indicating the different locations where proximity displays were tested.

Condition	Figure 6.7(a) (all albums)		Figure 6.7(b) (opened album)	
	Position	Display	Position	Display
Audit	2 (sidebar)		5 (sidebar)	
Facebook	3 (icon)		6 (icon)	
Mixed	1 (under)		5 (sidebar)	
Sidebar	2 (sidebar)		5 (sidebar)	
Under	1 (under)		4 (under)	

Table 6.3: Position and type of access-control proximity display shown for each condition and page.

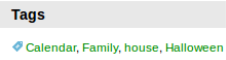




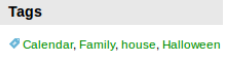
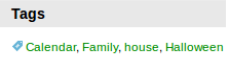
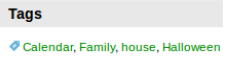


Condition	Figure 6.7(a) (all albums)		Figure 6.7(b) (opened album)	
	Position	Display	Position	Display
Audit	2 (sidebar)		5 (sidebar)	
Facebook	3 (icon)		6 (icon)	
Mixed	1 (under)		5 (sidebar)	
Sidebar	2 (sidebar)		5 (sidebar)	
Under	1 (under)		4 (under)	

Table 6.4: Position and type of tag proximity display shown for each condition and page.

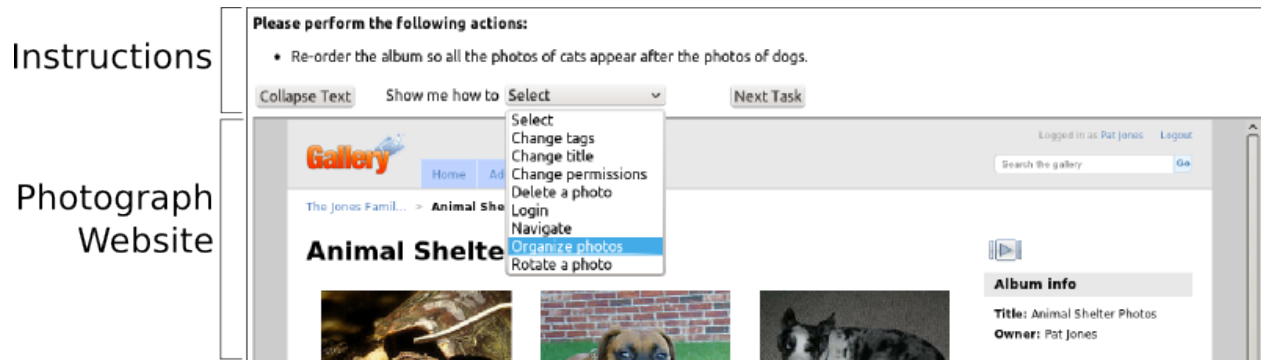


Figure 6.8: Screenshot from the online study showing the instructions and website frames. A control bar at the bottom of the instruction frame allowed participants to shrink the frame, obtain instruction on Gallery 3's features, and move to the next task.

In each task your friend or family member will ask you to perform a set of actions similar to the training. In addition to these actions you should also make sure the following statements are true:

- No spelling errors.
- Albums are tagged with the name of the friend or family member who took the photos.
- Photos are not sideways.
- Family albums can only be viewed by Family, and friend albums can only be viewed by Friends.
- No blurry photos.
- Pat can view, add, and edit all albums.

(a) Personal scenario ideal policy

In each task co-workers, who work with Starlight Phones and Purse Central, will ask you to perform a set of actions similar to the training. In addition to these actions you should also make sure the following statements are true:

- No spelling errors.
- Albums are tagged with the name of the co-worker who took the photos.
- Photos are not sideways.
- Starlight Phone's albums can only be viewed by contractors from Starlight Phone and Purse Central's albums can only be viewed by contractors from Purse Central.
- No blurry photos.
- Dezig Design co-workers can view, add, and edit all albums.

(b) Work scenario ideal policy

Figure 6.9: Ideal policy rules in the online study.

Study group	Completed study	Changed permissions	Dropped out
sidebar	139	114	33
under	131	116	39
mixed	124	114	44
facebook	128	112	43
audit	136	126	32

Table 6.5: The number of users in the online study who completed the study in each condition, the number of participants who made at least one change to permissions in either condition, and the number of participants who agreed to the consent form but did not complete the study.

giving participants a second chance to identify and correct the errors. When participants started task 8, multiple errors, including permissions and tags, are introduced into the album by a script. The task email indicates that errors have been introduced but does not specify what the errors are or how many errors there are. Earlier studies showed that participants do not expect album content to change on such short notice, unless warned. Many participants simply did not believe that errors could be introduced so quickly or that the person emailing them would be that careless.

Task	Permission Error	Tag Error
1	Wrong group can view	Missing
2	No error	No error
3	Everybody on Internet can view	Missing
4	No error	No error
5	Same as task 3	Same as task 3
6	Extra group can view	Wrong person
7	No error	No error
8	Wrong group can view	Missing

Table 6.6: Tasks and their associated permission and tag errors.

After participants completed the training and tasks for both conditions, they were asked to fill out an online survey. The survey asked them to recall permissions and tags from both Gallery 3 sites. To test whether participants were aware of the ideal policy, we asked them questions about the permissions both sites should have had. These questions were shown to participants in random order to prevent ordering bias. Participants were also asked about past negative experiences, their own impressions about how many errors they found, and demographics.

6.4.4 Data analysis

All data from the study was placed in a MySQL database for analysis. During the study a log was kept of all the actions (change cover, delete, move, permissions, rename, and rotate)

that participants engaged in, as well as all the changes to photos and meta-data participants made. After each participant completed the study, scripts automatically collected, graded, and archived all data from the session. Permission and tag information was collected, compared with the correct permissions, and compared with the default permissions. After completing the study, participants used Survey Gizmo to fill out the end survey. This data was downloaded after all participants were finished and the data was put into the MySQL database.

Permission and tag correction

In this study we were unable to observe each participant's behaviors and were therefore unable to determine when they checked the permissions using the proximity display. Instead, for this study, we measured whether participants corrected the permissions/tags, as this is a strict subset of the number of participants who checked for a permission/tag error. In prior studies we observed that some participants easily internalize that permissions are important and feel inclined to change them, while other participants are unlikely to do so. In this study we observed that 37% of participants never made any change, correct or not, to the permission settings on either condition unless explicitly instructed to do so. Additionally, all but six of those participants also never made any change to tags without explicit instruction.

Participants interacted with five tasks that had permission errors. Task 5 was, from a permission error standpoint, a second chance to correct the permission error in task 3, so we considered these two tasks together and only evaluated the permissions at the end of task 5. For each task we compared the resulting permissions to the default ones and marked them as either *correct* or *wrong*. We also recorded whether permissions were changed but were still inaccurate. We then summed up the number of tasks where permissions were correctly changed for each condition participants saw. A Shapiro-Wilk test for normality showed the data to be non-normal, so we used a Wilcoxon signed rank test to compare participants' permission and tag correcting performance on the control and experimental conditions. The permission corrections were part of the planned tests, and the remaining statistical tests were corrected using the Holm-Bonferroni method.

Non-permission error correction

Gallery 3 stores the meta-data of all photos and albums in a MySQL database. After each study session several scripts were run against this database to collect and code relevant information and then archive the database. The scripts compared the meta-data on all albums and photos to the default meta-data. If the default value was different from the final value, then the album was marked as having been rotated, re-titled, deleted, cover changed, or permissions changed depending on the meta-data. A second script was run that coded each meta-data element as *error*, *error fixed*, or *no default error*. For example, some tasks required participants to make any change to the title of a photo. For these tasks any change to the title was considered correct. Some tasks required participants to make a specific change; in these cases, the applicable meta-data value had to match the

10. For the Family Calendar album which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Family Calendar Album	_ can currently view Family CalendarAlbum
Everybody on the internet	<input type="checkbox"/>	<input type="checkbox"/>
Adventure Friends	<input type="checkbox"/>	<input type="checkbox"/>
Animal Shelter	<input type="checkbox"/>	<input type="checkbox"/>
Family	<input type="checkbox"/>	<input type="checkbox"/>
Pat Jones	<input type="checkbox"/>	<input type="checkbox"/>
I don't know	<input type="checkbox"/>	<input type="checkbox"/>

Figure 6.10: Sample permission recall question from the post-survey. The question asks participants to recall both what the permissions should have been and what the permissions were at the end of the study.

required final state.

Memory

As part of the post-study survey, participants were asked recall questions about six albums, five user groups, and one action (view). Participants were also asked what the permissions should have been according to the ideal policy. A sample question can be seen in Figure 6.10. The order of recall and ideal policy application questions was randomized for each participant.

Participants were asked about the album they used in the tutorial, the album that changed permissions (tasks 3, 5, and 8), and the album from task 6. We decided to use the album from the tutorial because all participants would have seen and modified the permissions on the album. The other two albums were selected because one of them is used on multiple tasks, and the other is used on only one task. Additionally, all three albums have incorrect permissions by default so participants should only be able to answer correctly if they have read the permissions. While analyzing data, we discovered a data collection error where we mistakenly asked participants about the album associated with task 7 instead of task 6 when asking about the work scenario.

At the end of the study, we archived the final state of the permissions for all users. We downloaded the answers to all survey questions from Survey Gizmo and loaded them into the database. We then compared the final permission state to what participants thought were the correct permissions. The result was coded as “correct,” “wrong,” or “I don’t know.” The answers to the ideal policy comprehension questions were compared to the correct settings for those albums and were also coded as “correct,” “wrong,” or “I don’t know.”

6.5 Conclusion

In this chapter we presented the details of the methodologies used in the eye-tracker, lab, and online studies. All three studies were designed to test the same set of hypotheses, but used different methodologies and collected different types of data. The next chapter presents the combined results from all three studies.

Chapter 7

Effectiveness of proximity displays

In this chapter we present the findings from the eye-tracker study, lab study, and online study. We empirically test whether different styles of proximity displays improve participants' ability to identify policy errors and recall the implemented policy when compared to interfaces in which the implemented policy is shown on a secondary page. We also examine how people interact with permissions when they are located on a secondary interface and how showing them on a proximity display changes participant behavior.

Proximity displays are intended to help end users identify permission errors and improve users' understanding of their implemented policy, without negatively impacting their ability to complete their primary tasks. Theoretically, proximity displays do this by making it easy for people to notice the implemented policy during their normal interactions with their online photo albums.

In the eye-tracker study, we found that people looked at proximity displays located under photo/album thumbnails throughout the task, but tended to change permissions at the end of tasks, rarely interacting with the displays in the middle. In the lab study we found that participants were able to glance at proximity displays and quickly determine whether a policy error existed. We also observed that some participants benefited more from proximity displays than other participants. We ultimately showed, in the online study, that some proximity display designs do positively impact people's ability to identify and fix permission errors. Additionally, none of the designs negatively impacted participants' ability to fix other errors. We also observed behaviors in our online study that support the anecdotal observations from our lab study.

We present the results of the eye-tracker, lab, and online studies in one chapter to give the reader a more holistic understanding of how users interact with proximity displays, and the effect showing these displays has on user behavior. To assist the reader in recalling the details of each study, Table 7.1 briefly summarizes the studies and Table 7.2 lists the conditions tested.

In the remainder of this chapter we first test our three main hypotheses (Section 7.1), then we qualitatively examine how people notice and correct policy errors (Section 7.2), followed by a discussion of each of the proximity-display designs we tested (Section 7.3), finally we discuss the limitations of these results (Section 7.4), and conclude the chapter (Section 7.5).

Name	Location	Type	Length	Tasks	Participants	Conditions
Pre-study	Lab	Between-subjects	1 hour	9	26	3
Eye tracker	Lab	Between-subjects	1.5 hours	12	34	3
Lab	Lab	Between-subjects	1.5 hours	16	33	4
Online	Online	Within-subjects	1 hour	16	658	5

Table 7.1: Methodologies used in each study.

	Control		Proximity display					Permission Modification	
Study	No Info	Tag Info	Under Photo	Sidebar	Mixed	Facebook	Audit	Full	Dialog
Pre-study	X		X	X				X	
Eye tracker	X		X	X				X	
Lab	X		X					X	X
Online		X	X	X	X	X	X		X

Table 7.2: The conditions tested in each study; details on each condition can be found in Section 6.4.1.

7.1 Hypothesis testing

In the online study, we empirically tested our three main hypotheses in a photo-sharing environment where security was a secondary task:

- H1: Correcting/checking permissions** Users who see permission information on a proximity display check and correct permission errors more often than users who see permission information on a secondary page.
- H2: Permission recall** Participants who see permission information on proximity displays can recall those permissions better than participants who see permission information only if they click to a second page.
- H3: Negative effects** Participants who see proximity displays take no more time, and correct no fewer non-permission errors, than participants who see permission information on a secondary page.

The online study had the largest number of participants and was designed to quantitatively evaluate each of five proximity-display designs. Recall that the online study was a within-subjects study, so each participant saw both control and experimental conditions. There were five control conditions to match the five proximity-display conditions; we refer to each pair of control and experimental conditions as a *treatment*. Details about the control and proximity displays we tested can be found in Section 6.4.1.

The statistical tests used in Tables 7.3, 7.4, 7.5, 7.6 are Wilcoxon signed-rank tests. A Shapiro-Wilk test on the number of permissions corrected, as well as permission recall, in each condition showed the data did not have a normal distribution. We used the Wilcoxon test instead of a t-test because our data was non-parametric. We use the Wilcoxon signed-rank test because our study was within subjects. All p-values shown, except those for planned tests, which are shown in Table 7.3, have been corrected using the Holm-Bonferroni

method.

7.1.1 H1: Correcting/checking permissions

We found in the online study that placing a proximity display with implemented-policy information under album and photo thumbnails (Wilcoxon, $p=0.045$), or under album thumbnails and on the sidebar (Wilcoxon, $p=0.023$), resulted in participants correcting statistically significantly more access-control permissions than they corrected in the respective control conditions. However, there was no statistically significant difference in the number of permissions corrected between the control and experimental conditions using the sidebar treatment (Wilcoxon, $p=0.052$), the treatment that emulated Facebook’s proximity icons (Wilcoxon, $p=1.0$), or the treatment that showed information about who had previously viewed the album (Wilcoxon, $p=0.953$).

To better understand the difference between how proximity displays impact the way people interact with permissions, as opposed to with other types of settings, the control conditions used proximity displays to show keywords that albums were tagged with, rather than permission information. We saw no statistically significant difference in the number of tags corrected when showing tag information on a proximity display (control) and when showing permission information on the proximity display (experimental). The largest difference between control and experimental was observed in the under-photo treatment where participants corrected an average of 0.91 *fewer* tag errors if they saw the tag information on a proximity display (control) than on a secondary interface. In both the under-photo and mixed conditions, participants were more likely to correct tag errors if they saw permission information on the proximity displays than if they saw tag information. This is a surprising result, we would expect that participants would correct more tag errors when they see tag information on the proximity display.

Condition	Permissions corrected out of 4						
	Wilcoxon p-value	Control			Permissions on proximity		
		Median	Average	StDev	Median	Average	StDev
under	0.045	0	0.924	1.316	1	1.176	1.444
sidebar	0.052	0	0.784	1.19	0	1.007	1.283
facebook	1	0	1.094	1.422	0	1.094	1.45
mixed	0.023	0	0.774	1.202	0	1.048	1.378
audit	0.953	1	1.14	1.394	0	1.147	1.443

Table 7.3: Results of Wilcoxon signed-rank statistical test (within-subjects). We present both the median and the average number of permissions corrected in the control and experimental conditions.

7.1.2 H2: Permission recall

After completing all tasks in both conditions, the online study participants were asked to fill out a survey. They were also asked to recall the current permissions for three albums

in each condition, all of which initially had permission errors. Participants answered 30 permission recall questions about three albums, all five user groups, and one action (view). In the analysis of permission recall, we exclude results from the training albums, on which all participants were forced to edit permissions. This leaves four albums: two from each condition. A participant’s answers were compared to the permission settings at the end of the study. Additional details about the data analysis and question format can be found in Section 6.4.4.

The results from the two non-training albums per condition showed no statistically significant difference in the number of permissions recalled between conditions in any of the treatments (Table 7.4). Participants recalled an average of 6.5 of 10 permissions. For comparison, participants recalled an average of 7.5 permissions out of 10 for the training albums.

Condition	Permission settings recalled out of 10						
	Wilcoxon signed rank	Median	Average	StDev	Median	Average	StDev
under	1.0	7	6.466	2.813	8	6.779	2.946
sidebar	1.0	7	6.345	2.807	8	6.525	2.793
facebook	1.0	7	6.586	2.912	7	6.414	2.822
mixed	1.0	7	6.242	2.73	7.5	6.742	2.814
audit	1.0	7	6.676	2.751	7	6.721	2.685

Table 7.4: The online-study participants’ ability to recall permission settings for two non-training albums (five questions each). Reported p-values reflect Holm-Bonferroni correction.

Participants in the mixed treatment showed the most improvement in recall between control and experimental conditions, however, the user group (e.g., Family, Friends) that participants were asked about also had a large effect on recall. Recall was highest on questions about groups that never had permission errors associated with them (Pat and Dezig). In the experimental condition, 81.2% of participants accurately recalled permissions for these groups, compared with 79.9% of participants in the control condition. These groups never had an error associated with them, so their initial, ideal, and final states should be identical and participants had no reason to change them. The worst recall results were associated with the training user group and the user group associated with the album that changed. Participants recalled between 46% and 51% of permissions associated with these groups, and the percentage did not vary significantly by condition. Recall questions about permissions for these groups were also the most likely to be incorrectly answered (participants could also indicate “I don’t know”). Participants answered between 25% and 38% of these questions incorrectly.

In addition to recall questions, we asked participants what the correct permissions were for these albums. No treatment exhibited a statistically significant difference between conditions (Table 7.5). This was expected and shows that participants in all treatments understood the correct permission state for each album.

	Ideal policy recalled out of 10						
	Wilcoxon	Control			Proximity		
Condition	p-value	Median	Average	StDev	Median	Average	StDev
under	1.0	8	6.809	3.55	9	7.099	3.601
sidebar	1.0	8	6.813	3.564	9	7.108	3.629
facebook	1.0	9	7.25	3.514	8	7.133	3.249
mixed	1.0	8	6.855	3.414	9	7.113	3.45
audit	1.0	8	7.125	3.332	9	7.199	3.326

Table 7.5: The online-study participants’ ability to apply the permission rules in the ideal policy for the two non-training albums per condition (5 questions each). Participants were asked what permissions Pat/Pat’s boss would have wanted to set. Reported p-values reflect Holm-Bonferroni correction.

7.1.3 H3: Negative effects

As can be seen in Table 7.6, participants in the online study exhibited no statistically significant difference in the number of non-permission tasks corrected. Because the online study was time limited, we cannot make any claims about the time required to complete the tasks. In the lab study, which was not time limited, the control condition and experimental conditions showed no significant difference in either time to complete the tasks or number of non-permission errors corrected. We therefore conclude that proximity displays do not negatively impact time or accuracy of other tasks.

	Non-permission and non-tag errors corrected out of 37						
	Wilcoxon	Control			Proximity		
Condition	p-value	Median	Average	StDev	Median	Average	StDev
under	1.0	27	26.588	4.474	27	26.672	4.657
sidebar	1.0	28	27.079	4.188	27	26.698	4.995
facebook	1.0	28	27.68	3.669	27	27.102	4.819
mixed	1.0	27	26.75	4.121	27	26.823	3.984
audit	1.0	27	27.353	4.076	27	26.882	5.369

Table 7.6: In addition to tag and permission errors, the online-study participants were asked to correct issues with the titles, organization, orientation, and content of photos. This table reports the number of non-permission and non-tag errors participants corrected out of 37 errors. Reported p-values reflect Holm-Bonferroni correction.

7.2 How people notice and fix permission errors

The lab study was designed to collect a large amount of qualitative data to better understand how participants notice permission errors. This section focuses on results from the lab study; where appropriate, we also present data from the online study that supports

or contradicts our lab study conclusions, and from the eye-tracker study to discuss when participants look at proximity displays.

We found, using the eye tracker, that participants saw permissions on proximity displays throughout each task (Section 7.2.1) but they appeared to correct the permission errors at the beginning and end of tasks (Section 7.2.3). We observed that some participants checked permissions rarely; these participants benefited the most from seeing proximity displays. Conversely, some participants checked permissions frequently with little provocation. These participants tended to check permissions less often when shown a proximity display (Section 7.2.2).

7.2.1 Noticing permissions

One of the goals of proximity displays is to enable participants to notice and check permissions quickly and easily. In the lab study we observed that participants who saw proximity displays seem to check permissions less often than control participants, but both groups corrected the same number of permissions. In this section, we explore how people become aware of an error in their permissions, specifically looking at the process people go through when finding permission errors using the displays.

To understand how people are noticing permission errors, we use data from the eye-tracker study, where we used an eye tracker to determine when participants were looking at proximity displays. Figure 7.1 is a histogram of the number of instances when participants, in the under-photo and sidebar conditions, fixated on a proximity display. A *fixation* is an eye-tracking term for when participants' gaze rests on a single point on the interface. We only counted fixations on the webpage participants worked with in advance of modifying permissions, and we normalized fixation times to make them comparable across participants. The majority of participants stayed on a single album page for the length of a task. Time 0 on the graph in Figure 7.1 represents the instant when participants opened the album page, and time 100 represents the instant when participants opened the permission-modification interface that is located on a new page. Participants in all three conditions spent an average of 4.4 minutes on a page before opening the permission-modification interface. We observed that under-photo participants fixated on the proximity displays throughout the task, but explicitly checked and corrected permissions at the end of the task (Figure 7.1(a)). Sidebar participants looked at the display just before transitioning to the permission-modification interface, but looked at the display rarely before that point (Figure 7.1(b)). Participants tended to read the printed email at the beginning and end of tasks, so the slight decrease in the number of fixations at the beginning and near the end is likely caused participants not looking at the screen, rather than participants choosing not to look at displays during those times.

Given that the under-photo condition places proximity displays all over the screen, it is difficult for participants not to see or fixate on a display during the course of a task. Therefore we have to ask whether participants are paying attention to the permissions, or just looking at the proximity displays without absorbing the information they are showing. To test whether participants were paying attention to the displays before changing permissions, we designed the lab study to have permission errors randomly

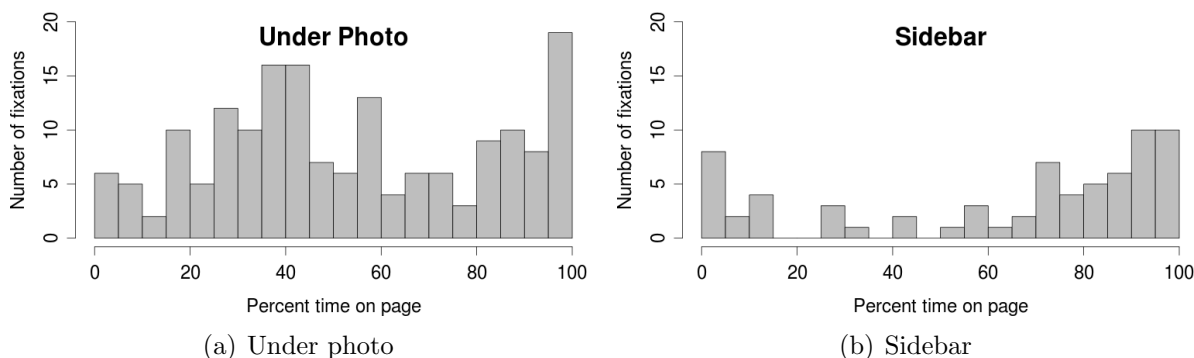


Figure 7.1: Histogram of the number of fixations on proximity displays for all participants (y-axis) against the amount of time spent on the page, normalized (x-axis). Data is from the eye-tracker study.

ordered and assigned to tasks. This minimized the effect task wording and order had on when permissions were checked. Further description of the lab-study methodology can be found in Section 6.3.2.

To better understand whether participants were able to notice permission errors easily, we needed to measure noticing behavior separately from permission-correcting behavior. We use the term *noticed* when discussing what participants actually saw; this is what we are attempting to detect and measure. In the lab study we approximated when participants noticed permission errors by measuring *explicit checking* behavior. Control participants were said to have *explicitly checked* permissions if they opened the permission-modification interface. Participants who were shown proximity displays were said to have *explicitly checked* permissions if they (1) opened the permission-modification interface; or (2) read the permission aloud; or (3) indicated through mouse behavior that they were reading the permission display (moving the mouse under the words while reading or circling the display with the mouse); or (4) pointed at the proximity display with their hand while staring intently at the screen. The identification of explicit permission-checking behavior in the under-photo condition was done by the researcher during the lab study.

We compared the number of times participants explicitly checked permissions on tasks with errors to tasks without errors (Figure 7.2(a)). Participants in the control condition were equally likely to explicitly check permissions on tasks with and without permission errors. This is expected, because participants had no way of knowing whether an error existed without opening the permission-modification interface. Consequently, for the control condition we had a very accurate measurement of how often a permission was checked. Participants in the under-photo condition were more likely to explicitly check permissions on tasks with permission errors than on tasks without permission errors. On average, under-photo participants explicitly checked permissions on 3.2 tasks with errors and 1.7 tasks without errors. This suggests that participants are paying attention to the display more often than we are observing through measuring explicit-checking behavior, because our measurement definition did not capture all the permission-checking events.

Why, in the lab study, are we not observing every permission-checking event? The

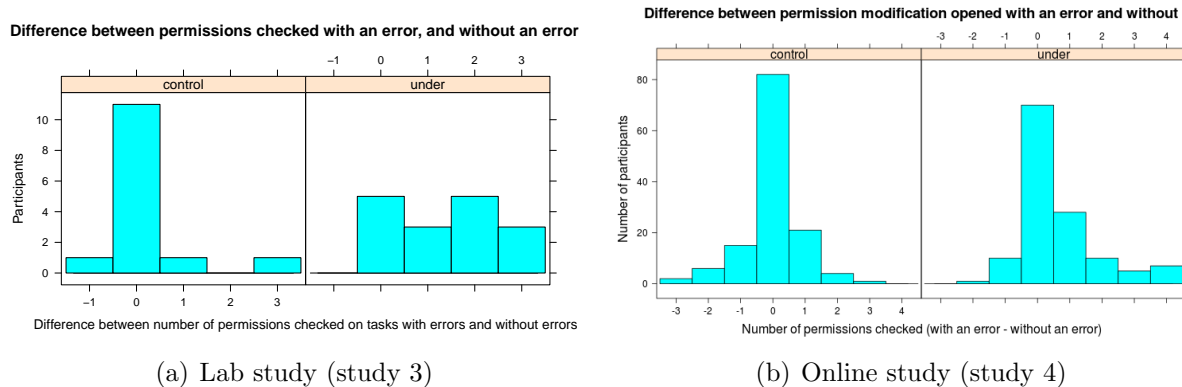


Figure 7.2: Histogram of the difference between the number of tasks on which participants checked permissions when the tasks had an error and the number of tasks on which they checked when the tasks did not have an error. For example, in the lab study 11 control-condition participants checked the same number of permissions in tasks with errors as they did in tasks without errors. In the lab study (Subfigure (a)) “checked” was defined based on behavior observed by the researcher administering the study. In the online study (Subfigure (b)) “checked” was defined as opening the permission-modification interface.

eye-tracker data suggests that participants are fixating on the displays midway through the tasks, but the permission-modification behavior suggests that they explicitly check permissions occasionally at the beginning and frequently at the end of the tasks. While participants’ eyes may be fixating on a proximity display, we cannot be sure that they are absorbing the information, and so the eye tracker is likely overestimating the number of times participants paid attention to a proximity display. Similarly, think-aloud data only captures information participants process sufficiently to articulate. We know from the lab study (Figure 7.2(a)) that participants are identifying the absence of an error and not indicating this verbally or through other behavior that we measure. Think-aloud theory tells us that this is because the information is either: (1) non-linear and therefore challenging to verbalize, or (2) in working memory for a very short time [100].

When asked, lab-study participants reported noticing permissions while working on other issues. Because they were distracted, they put off fixing the permissions until the end of the task, by which point they may have forgotten that they had noticed an error.

One participant in the under-full condition said:

I was more just focused on getting this done first. I felt like if I looked at the permissions like or if I glossed it over, I just wanted to get this stuff done first, and thinking that I would go back to it but I never ended up doing that.

If we look at participant behavior in terms of the C-HIP behavioral model [108] and the HITL framework [26] described in Section 2.5, we can better understand what is happening. The C-HIP model describes the set of states users must go through between when a warning becomes visible and the warning actually effecting behavior. According to the model, once users have looked at the warning (Attention Switch) they decide whether the warning is worth focusing on (Attention Maintenance). If the warning is worth focusing on, they try

to understand the warning and comprehend what the warning is saying (Comprehension and Memory). Once they comprehend the warning, they decide based on their beliefs whether the warning applies to them (Attitudes and Beliefs). Finally, if they consider the warning relevant, they may be motivated to do something (Motivation) that may result in a change in their behavior (Behavior). We hypothesized that our participants were making the attention switch (the eye tracker in the eye-tracker study reports fixations), and then glancing at the proximity display to see whether it “looked wrong” (Attention Maintenance). If participants believed the display was inconsistent with their expectations they would focus on it (Comprehension and Memory); that is when we observed them explicitly checking permissions. At this point, our participants either made the permission change or waited until the end of the task and then made the change (Motivation and Behavior). However, if the display appeared to match their expectations, participants saw no need to continue focusing on the display (Attention and Maintenance), and moved on to other tasks before they tried to comprehend the content (Comprehension and Memory).

We hypothesize that because we gave participants a single access-control policy that was globally applicable, they were able to learn to quickly differentiate between proximity displays that showed correct permissions and those that showed policy errors. This allowed them to glance at the displays and determine whether they looked correct during the Attention Maintenance stage of the C-HIP Model without having to move to the Comprehension Memory stage, where they would have 1) transformed the contents of working memory to a form that was easy to vocalize, or 2) had to keep the information in working memory for long enough to vocalize.

During the unstructured interview at the end of the lab study, the researcher asked participants how they had identified permission errors. Participants in the under-photo conditions talked about the heuristics they used. Instead of reading the whole policy, an under-full condition participant said he would just look to see whether Everybody could view the album:

If it was company related then it should say Everybody and if it didn't say Everybody then it was wrong, and I would know that just by looking.

Another under-full condition participant showed the researcher what correct and wrong policies looked like. His primary metric appeared to be the length and shape of the words on the display:

[Indicates proximity display on the screen] if there is like a lot of things I will look because there is only one a few things you should have up as the permissions.

In the online study, we measured the number of times participants opened the permission-modification interface (Figure 7.2(b)). We observed that under-photo, sidebar, and mixed condition participants were statistically more likely to open the permission-modification interface if there was a permission error than if there was no error ($p < 0.001$). Participants in the control conditions were equally likely to open the permission-modification interface regardless of whether a permission error existed. This shows that participants in the under-photo, sidebar, and mixed conditions were able to use proximity displays to

identify errors and avoid needlessly opening the permission-modification interface when an error did not exist.

These results tell us that proximity displays should provide sufficient information for users to determine whether an error likely exists without interacting with the display.

Checking permissions all at once

Another reason we decided to provide participants with a single ideal policy in the lab study was to provide a more natural environment. People typically know their own policies. By providing a single policy (rather than different policies for different albums), we allowed participants to choose when to make changes instead of forcing them to make changes during a specific task. We found that some participants, regardless of condition or permission-modification interface, have a tendency to take a single pass through the whole policy, correct all the permission errors, and never look at the permissions again. We term this behavior *checking all at once*.

Participants who check permissions all at once never checked any permission again after doing so, even when some of the decisions they made in their permission-correcting pass were wrong. As part of the unstructured interview, the researcher asked these participants whether they had considered looking at the permissions again, or if they had been concerned about errors. Some participants reported briefly debating whether a permission was wrong but decided to do nothing about it.

This observation tells us that some participants want to correct permissions in a single pass and not think about them again. This tendency appears to be independent of the treatment and condition, though some display designs might encourage the behavior more than others. The disinclination to check permissions again indicates that though these participants correct more permissions in their one pass, they are susceptible to missing errors introduced after the checking took place.

The online study was time limited, which discouraged participants from checking all at once. However, we still observed 92 participants (14%), from all treatments, who corrected permissions on more than one album during at least one task. Twenty of those users corrected permissions on more than two albums during a single task (3% of all participants).

7.2.2 Participants' tendency to check permissions

We observed a high variance between lab-study participants in their behavior towards permissions. Some participants completely ignored permissions, and some participants took checking permissions very seriously. This made it difficult to determine whether the permissions were being checked because of the interface or because participants were predisposed to check them. To account for the difference, we made the online study a within-subjects study so that each participant would experience both a control condition and an experimental condition.

The lab-study participants in both under-photo dialog and control-dialog conditions appear to check permissions either frequently or not at all (Figure 7.3(a)). We hypothesize that some unobserved variable causes some participants to frequently check permissions and

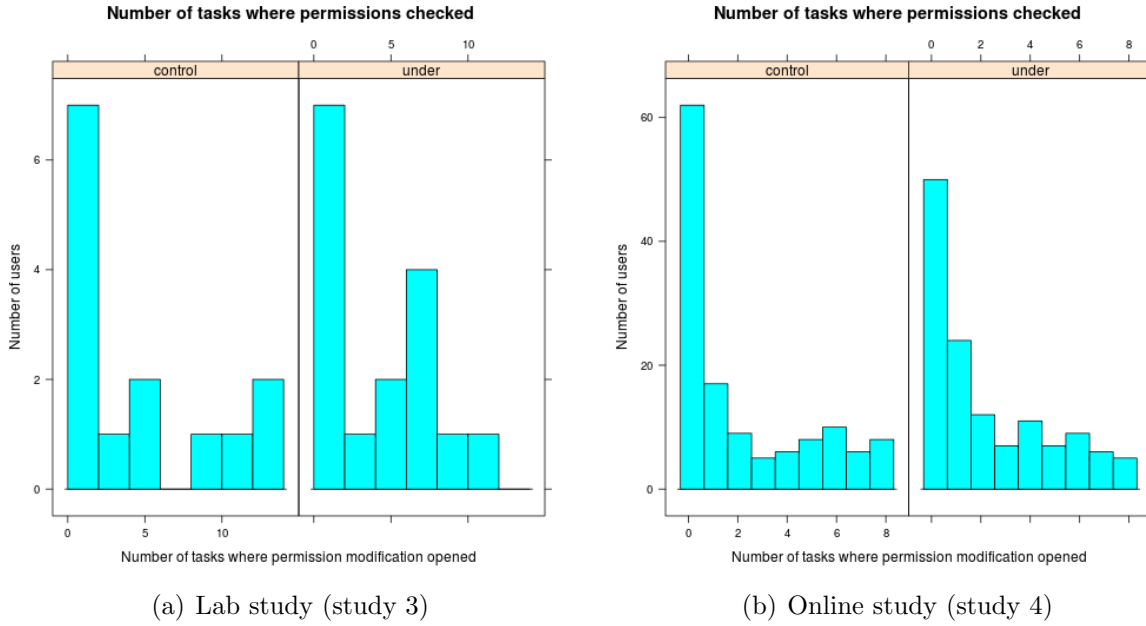


Figure 7.3: Number of tasks where the permission-modification interface was opened by participants in the control and under-photo conditions. Graph *a* shows data from the lab study, and graph *b* shows data from the online study. Graphs of other conditions in the online study are nearly identical.

some participants to rarely check permissions. We refer to these two groups as *infrequent permission checkers* and *frequent permission checkers*. In the lab study, we define frequent permission checkers as those who check permissions on more than half of tasks.

Infrequent permission checkers checked permissions on an average of 1.8 tasks, and frequent permission checkers explicitly checked permissions on an average of 8.3 tasks if they saw the proximity display (under), or 12.5 tasks if they did not see a proximity display (control). As discussed in Section 7.2.1, our measurement for explicit checking behavior underestimates the true number of times an under-photo participant checks permissions. Underestimating permission checking may be one of the reasons for the difference between under-photo and control participants who checked permissions frequently.

In the online study, we looked at the number of tasks where participants opened the permission-modification interface as a way to measure how often they were checking permissions. As previously mentioned, this is an accurate measure of the number of times participants check permissions during the control condition, but is a lower bound on the number of times experimental-condition participants checked permissions. Similar to the results from the lab study, many participants in the online study never checked permissions (Figure 7.3(b)). However, those participants who did check permissions tended to check a similar number in both conditions. There is a positive linear relationship between the number of tasks in which the permission-modification interface is opened in the control and experimental conditions regardless of the treatment (linear model, $p < 0.001$).

The decision to check permissions appeared to be an individual choice. We were not

able to determine any factor that caused participants to check permissions. In the lab-study post-interview, participants were asked why they did or did not check permissions. When asked why they checked permissions, most lab-study participants responded that checking permissions was part of their job, or that three of Gerald’s rules mentioned permissions (Figure 7.6). Participants who checked permissions in the majority of tasks could not even seem to understand why the researcher had asked the question. From their perspective, the fact that permissions should be checked was obvious. This was well summarized by a control-full participant who said:

But it said it was your job. You know what I mean, if you could lose your job because you screwed it up then why wouldn’t you

Some participants mentioned that managing permissions was important to them in their own lives, or that the rules resonated with them. One participant in the control-dialog condition said:

There were guidelines explicit in the instructions that had to do with what the boss Gerald wanted like access to certain albums. And also personally with my privacy settings on the Internet I want to make sure that my albums are only available to people I want.

Across conditions, 18 participants (54.5%) checked permissions in less than half of the tasks. When asked why they did not check permissions, participants came up with a wide range of answers, but the most common type of answer is well summarized by the following quote by a control-full condition participant:

I think I may have forgotten about the permissions.

During the unstructured part of the post-study interview, the researcher asked participants why they were not checking or changing the permissions despite the fact that they checked for and fixed non-permission errors. One participant in the control-dialog condition responded:

With all the other errors since they were right in front of me I could just see them and they kinda triggered my memory that way. I guess without the permissions error being there I couldn’t... it didn’t just pop in my head.

Despite the lab-study post-interview, we were not able to determine what variable caused some participants to forget to check the permissions and others to check permissions frequently. The post-survey in the online study included multiple questions to determine whether attitudes, opinions, or past experiences had an effect on participants’ tendency to check permissions during the control condition. We created these questions based on the work by Wang et al. on Facebook regrets [104], and Tsai et al. on online shopping privacy [99]. The only question to show any correlation was: “Do you agree or disagree with the following statement: Most businesses handle the personal information they collect about consumers in a proper and confidential way.” Participants answered using a five point Likert scale. Participants who agreed with this statement were more likely to change permissions in the control condition (corrected ANOVA, $p=0.043$).

Proximity displays impacted participants differently depending on their permission-checking behavior. In the lab study, participants who checked frequently appeared to be negatively impacted by the proximity display. The numbers were too small for meaningful

statistics, but the researcher observed several cases where an under-photo participant, who was clearly checking permissions on every task, forgot to check on one or two tasks. During the post-survey, however, they were certain that they had checked permissions on every task. We believe that by encouraging glancing at permissions, we decreased the amount of focus participants gave to the act of checking permissions and thereby increased the number of errors missed. Social psychology tells us that tasks that receive less focus are more likely to be forgotten [24]. We have seen this effect before in error-identification work [101], where participants felt that they would notice whether the screen showed an error, but failed to notice the error because it was insufficiently “obvious,” and were therefore very confident that no error existed.

Participants who checked infrequently showed the opposite trend: under-photo participants tended to check permissions on about one more album than control participants. This trend was visible in both the lab and online studies (Figure 7.3). Looking at the online study, we see that participants in the under-photo treatment corrected only 0.25 more permission errors in the experimental condition than the control condition, and participants in the mixed treatment showed only a 0.77 improvement on average (Table 7.3). These numbers reflect the fact that most participants either check permissions one or two more times in the experimental condition than in the control condition, or check permissions on neither.

Recall that proximity displays are intended to be passive and only be looked at by users occasionally. They assist infrequent permission checkers by enabling them to easily check for errors at any time. The result that participants in the under-photo and mixed conditions check a permission more often in the experimental condition than in the control condition shows that the displays are fulfilling their intended role of occasionally assisting people. Proximity displays may be less beneficial in environments where checking the permissions frequently is important, since participants may be more likely to miss errors when glancing than when explicitly checking.

7.2.3 When do people change permissions

One of the purposes of proximity-information displays is to display implemented policy to end users in a way that naturally fits into their normal work-flow. We want to show implemented-policy information to people at the time and place when they will most likely need it and be receptive to it. We wanted to know 1) where/when do people naturally become interested in permission information, and 2) how can we manipulate the proximity-display design to best support this?

When we began testing the effectiveness of proximity displays, we expected that the proximity displays’ spatial proximity to users’ main focus (the photos), would allow users to glance at the displays as part of their workflow. Hence, we expected that the under-photo condition would outperform the sidebar condition in which displays were not spatially located near participants’ primary focus point. The eye-tracker study showed this to be true, with the under-photo condition outperforming the sidebar condition [102], and later the online study showed the same thing (Table 7.3). However, putting the proximity display under every photo takes a large amount of screen real estate and potentially distracts users,

so we wanted to use our understanding of how permission errors are identified to find a more appropriate solution. We observed the following permission-checking behaviors: 1) participants check permissions at the beginning and end of tasks, 2) participants tend to view permission errors as dissimilar to the other errors they are looking for, and 3) when proximity displays are shown under album and photo thumbnails, participants tend to check permissions using the display located under the album thumbnail.

Participants check permissions at the beginning and the end of tasks

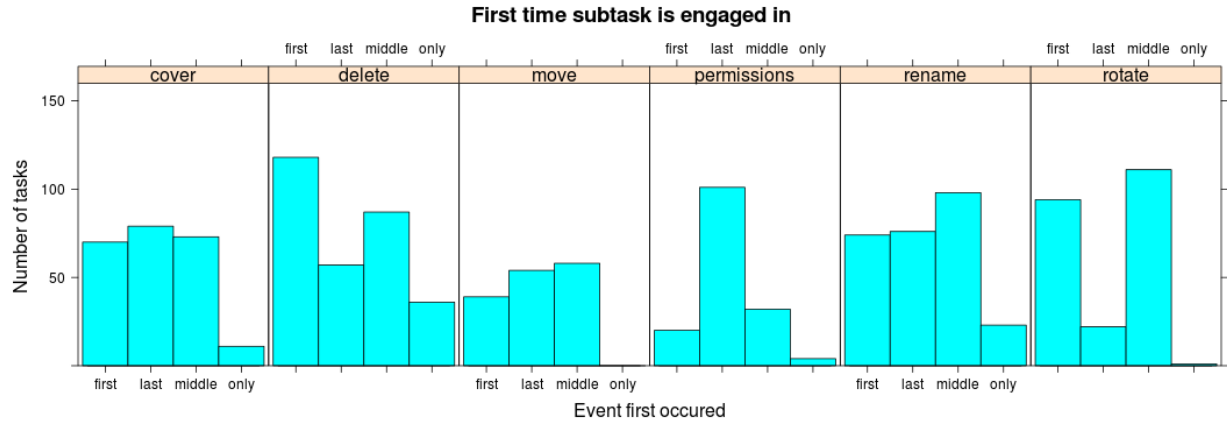
During the pre-study and eye-tracker study, we observed that the majority of the participants explicitly checked and changed permissions at the beginning and end of tasks. This was observed across conditions and across tasks. At the time, we hypothesized that the behavior was due to participants' need to go to a separate page in order to change the permissions. To test this hypothesis, we introduced a permission-modification-dialog condition into the lab study. Half of the participants were given the full-page permission-modification interface used in the eye-tracker study that required participants to switch webpages. The other half of the participants were given the permission-modification-dialog interface, which did not require switching pages. In this section we take a detailed look at when participants modified permissions as opposed to other types of actions.

Participants in the lab study engaged in a wide number of actions, and because they were not time limited, they took a wide range of times to complete tasks. Because of the wide variation in times, we analyzed the order participants chose to engage in the different actions instead of the time when the action was completed. The analysis of the order (Figure 7.4) showed that participants predominately engaged in a permission-modification action as the last action they engaged in. To determine this, we ordered the possible action types based on the first timestamp associated with each action type (which action was done first) for each user. We also ordered the possible action types based on the last timestamp associated with each action type (which action was finished last). We refer to the last time an action was engaged in as when the action was *completed*. A more detailed explanation of this analysis can be found in Section 6.3.4. There was no significant difference in action order between conditions.

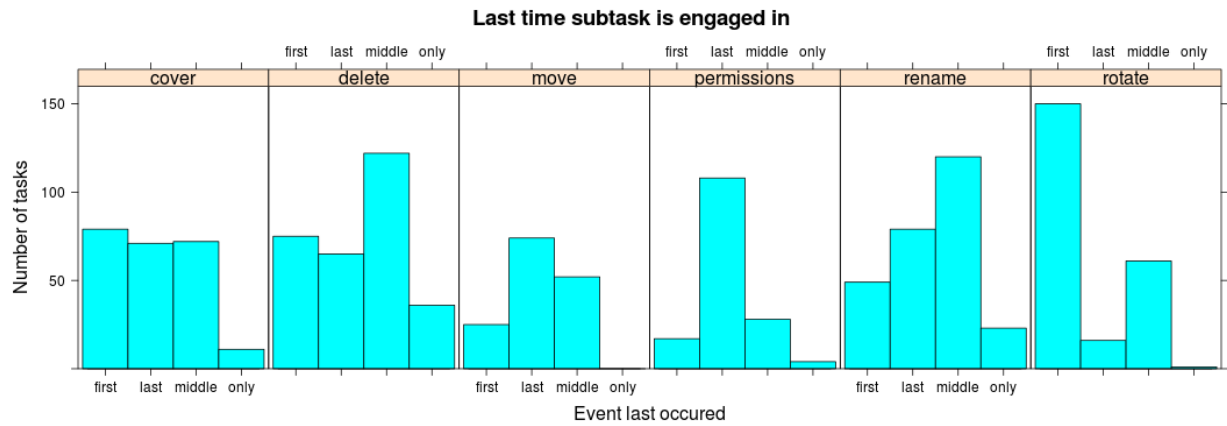
In the lab-study post-interview, participants who were obviously checking permissions at the end or beginning of tasks were asked why they were doing so. Most people did not know why they were checking at the end, and instead talked about how they had approached the tasks in general. The following quote from an under-dialog condition participant describes typical user behavior:

I think maybe because in the beginning I was jumping around just exploring the whole thing. And not really paying as much attention. Then I methodically went through and just, and it is pretty easy to just mouse over stuff, so it did not hurt to check.

Participants reported being very focused on the distractor task and the non-permission errors. They talked about changing permissions as a different type of action than the non-permission actions. One participant in the under-full condition explained:



(a) Event first occurred



(b) Event last occurred

Figure 7.4: While working on tasks in the lab study, participants were free to engage in actions in any order, including interleaving actions. For example: a participant could rotate a photo, delete a photo, then rotate a photo. Graph *a* shows the first time an action of that type was engaged in during a particular task and whether that action was the first action, the last, neither first nor last (middle), or the only action engaged in. The height of the bars indicates the total number of tasks across all users; the summation of all bars in a subgraph is the number of tasks, across all users, in which the action was engaged in at least once. Graph *b* is similar to graph *a* except that it shows the last time a action is engaged in during a task.

I guess I read the task and did everything that was required of me and left monitoring, personal monitoring for, you know, the last stuff. It was just easier to do what I had to do first, or just perform the request and then make sure that the policy was followed.... Seemed more intuitive the way I did it.

A control-full user explained this tendency to put permissions last the best. She likened setting the permissions to remembering to turn off the oven and then later decided it was closer to locking the door at night. Both were tasks that she always had to explicitly remember to do before going to bed.

Just because it comes at the end doesn't mean that it is unimportant to me. It probably means that ... That is like the final, this is it now. You do everything you are supposed to do before you go to bed then you make sure you lock the door. So that is like locking the door, checking those permissions, that is like the final security piece. – *Participant in the control-full condition.*

Participants deliberately decided not to modify permissions until they were finished checking the other requirements (Figure 7.4). Visually obvious actions such as rotating sideways photos are engaged in first or in the middle, and are the first action to be completed. Less visually obvious actions such as renaming, which includes both correcting spelling errors and making changes specified by the email, are initiated at any time and are rarely the first action to be completed. Actions that require a large dialog and focused attention, such as organizing photos (move), are engaged in at any point with a minor bias towards later and tend to be completed last.

After the lab study, we decided that we needed to better support permission-checking at the beginning and end of tasks. We were also concerned that putting information under every photo was causing participants to naturally ignore it due to habituation. Finally, we were concerned that when participants saw proximity displays everywhere, they just assumed that they would spot an error if it existed and therefore did not focus on the displays enough to actually check for errors.

To address these concerns, we added the *mixed* condition to the online study. This condition shows the proximity display under the album thumbnail, and when the album is opened the proximity display is shown on the sidebar instead of under every photo. The intention was to encourage participants to look at the permission information as the album was opened, or at the end of the task after the album was closed. Participants who choose to check permissions in the middle could easily do so with a quick glance at the sidebar, but the information was not near the photos participants were actively working with. As can be seen in Table 7.3, the mixed condition outperformed all other conditions in terms of number of permission errors corrected.

Permissions perceived differently from other errors

Participants talked about how challenging it was to have to think about finding both permission and non-permission errors at the same time. Participants were given five rules (Figure 7.6) that they were supposed to enforce when interacting with albums. However,

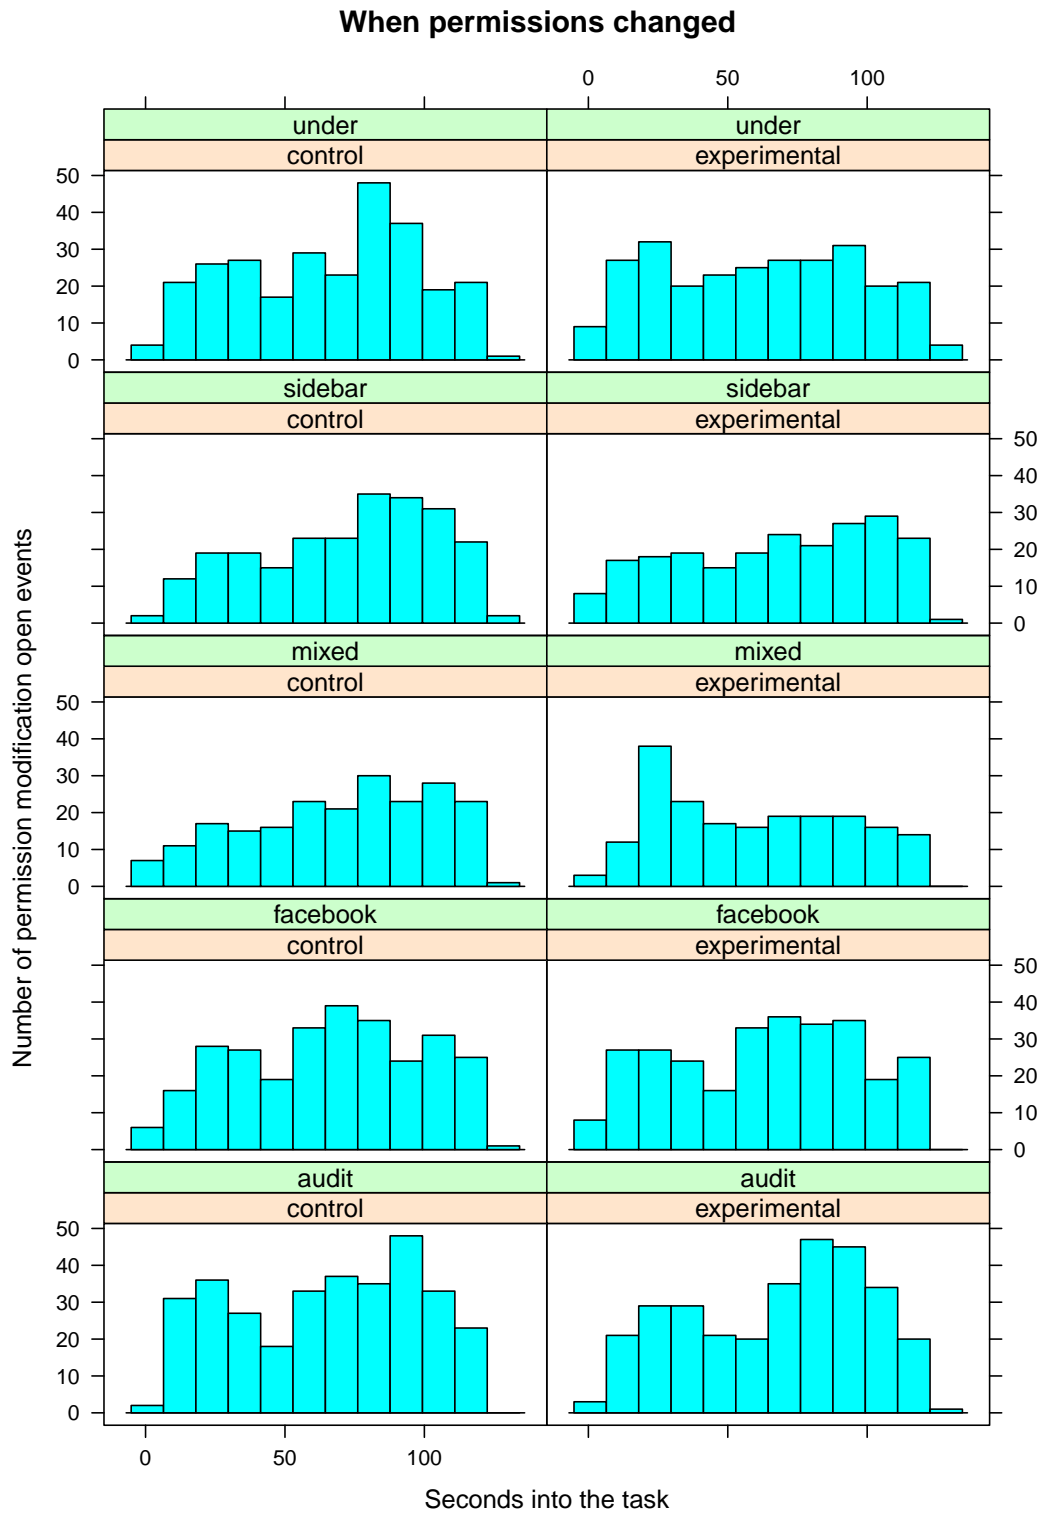


Figure 7.5: The number of seconds into the task when the permission-modification interface was opened by participants in each condition. Events from task 1 and the training are excluded to remove bias caused by prompting participants.

they appear to have considered the permission rules to be different from the others. One participant in the under-dialog condition explained:

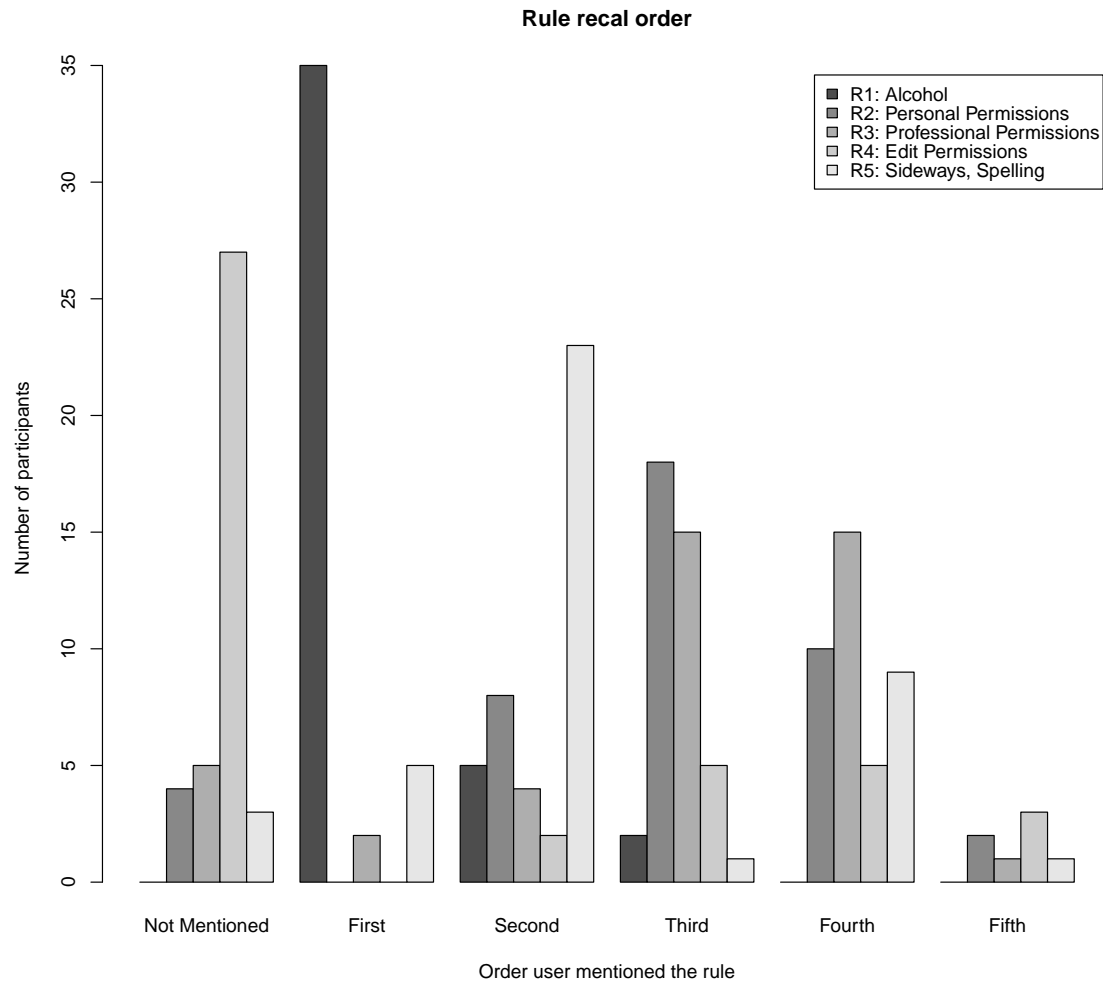
It was hard it was kinda balancing two aspects, it was either like maintaining the policy like the whole alcohol thing and at the same time making sure it was like a ... it is not just open to everybody it was exclusive to some people who can see it and understand it. If it was selective in that sense that people could see it that you wanted them to see it, you know, the alcohol might have been ok or the policy might not have applied as strictly. But it was like trying to balance.

As part of the lab-study post-interview, participants were asked to recall the boss's rules (ideal policy) in their own words. Figure 7.6 shows how many participants remembered the policy rules and the order in that they recalled the rules. The majority of participants first recalled rules 1 and 5, which have to do with alcohol, blurriness, rotations, and spelling errors; and then recalled rules 2 and 3, which concerned permissions. Rule 4, which specified who could add to or edit albums, was rarely recalled. There was no significant difference in the rules remembered among conditions.

Combined with post-study interviews, the information from Figure 7.6 suggests that participants are mentally grouping permission rules as different from the non-permission requirements. Additionally, participants appear to think of permissions after thinking about the other types of requirements. The way people group types of errors is important because it may help explain why participants change permissions first or last. If permissions are perceived as different than other attributes that could be manipulated, then checking them may require participants to swap out working memory. People would not want to change what they are thinking about multiple times in a task, so they wait until the end and change what they are thinking about then.

Participants in all treatments tended to open the permission-modification interface at the end of the task when they were in the control condition — recall that each participant was exposed to both a control condition and an experimental condition (Figure 7.5, column 1). However, participants who saw permission information in the mixed condition (Figure 7.5, column 2, row 3) did the opposite and tended to open the permission-modification interface at the beginning of the task. This is particularly notable since the same participants had the opposite behavior in the control condition.

Participants in the lab study talked about how permission errors were different than the other types of errors they were looking for. One potential difference between the types of errors might be participants' pre-study understanding of what "correct" and "wrong" states look like. Participants entered the study with a well-practiced ability to identify spelling errors and sideways photos. We did not have to impart what correct and wrong states were for such errors. Even the rule about no alcohol in photos was reasonably familiar to users, and several commented how they normally do not post that type of photo. However, participants had to be told what the "correct" and "wrong" permission settings were. Therefore, permission errors might have been different because participants did not have prior experience identifying those types of errors.



Gerald's Photograph Policy

1. No photographs containing drugs, alcohol, or anything inappropriate.
2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.
3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.
4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.
5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

Figure 7.6: As part of the lab-study verbal post-survey, participants were asked to recall Gerald's rules, in their own words. The above graph shows the order in which participants recalled the rules. The majority of participants recalled the rules in the following order: R1, R5, R2, R3 and forgot to mention R4.

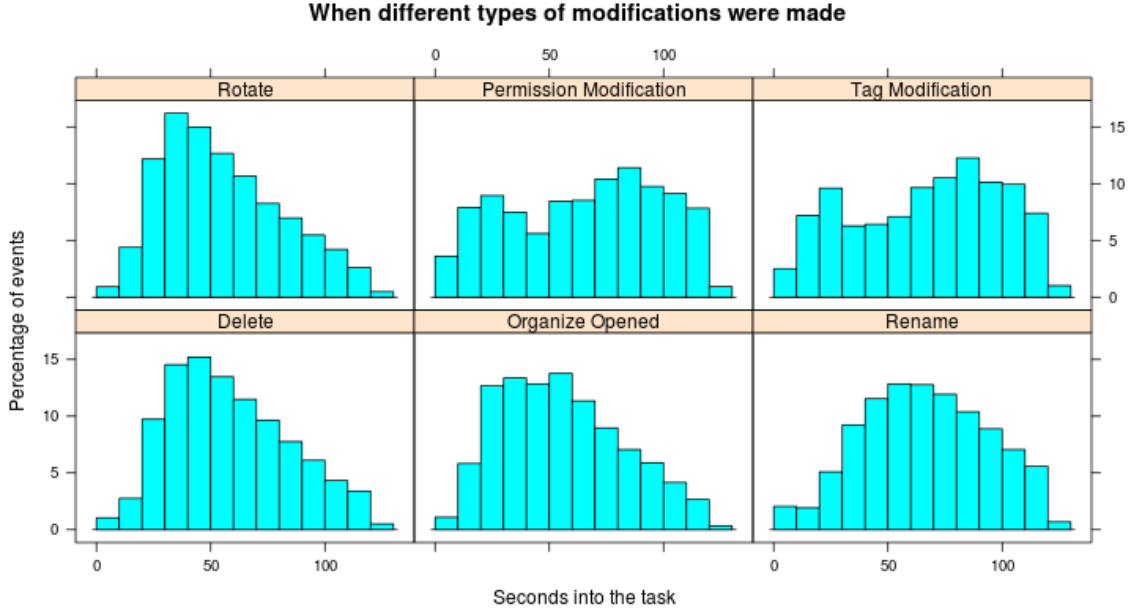


Figure 7.7: Number of seconds into a task that an action was engaged in (10 second intervals). Histograms show all participants across all conditions, both with and without permission proximity displays. Events from task 1 and the training are excluded to remove bias caused by prompting participants.

We wanted to know whether permissions are modified first and last because people intrinsically modify them then, or because participants had to learn to identify new “correct” and “wrong” states for this study. In the online study, we tried to account for this by introducing tags as a new action type that could have an error. As we did with permissions, we told participants what “correct” and “wrong” tag states looked like. Because the correct tag state was artificial, participants would have had no prior experience looking for these errors. We also put tag information on proximity displays in the control condition so participants received the same type and amount of exposure to tags as they did to permissions.

Figure 7.7 shows at what point during tasks participants engaged in each type of action. The distributions of times for rotate, delete, organize, and rename actions resemble skewed normal distributions, with the majority of participants making changes at similar times. Rotate and delete actions are typically performed first, while the more complex rename and organize actions were performed in the middle. These results are similar to the ones we observed in the lab study (Figure 7.4). However, if we look at the graphs for permission and tag modifications in Figure 7.7, we see that participants modify both tags and permissions at the beginning and end of tasks. This suggests that the reason participants are modifying permissions first and last is because the “correct” state is one participants have no prior experience identifying. Therefore, it may be that participants needed to focus more attention on the permissions than on the other actions to determine whether an error existed.

While familiarity with identifying correct and wrong permissions states may have been a factor in participants' tendency to check permissions last, it is unclear whether familiarity would significantly alter behavior. We observed these participants accurately identifying permission errors in advance of explicitly checking for them. Despite noticing permission errors quickly, the lab-study participants still modified permissions as the last action. The artificial ideal policy was likely a factor that influenced when participants changed permissions; however, we believe that how people think about permissions was also a factor.

Where permissions were checked

In the lab study, under-photo condition participants could check permissions in several places: under the album thumbnail, under any of the photo thumbnails, or by opening the permission-modification interface from the album thumbnail view or the photos view. The researcher noticed that some participants in the under-photo conditions would exit the album when finished with a task, mouse over the album thumbnail, explicitly check that the permissions were correct, and then declare themselves finished with the task. One under-full participant had this built into such a routine that he did not believe that the proximity display was even visible once an album was opened:

So once I open then I finish what I'm doing with the task and go back and look at the album and look at the permissions because you can't just see it right away. Um. Or you have to do it right away and then perform the task.

Conversely, some participants would click on the album thumbnail to open the album and appear to glance at the permissions while the album page was loading. These participants did not explicitly check permissions according to our definition, but the eye tracker indicates that they were looking at the display.

In the online study, we observed that the under-photo and sidebar participants tended to check permissions at the beginning of tasks in the experimental condition, as opposed to the control condition, in which they tended to check at the end of tasks. We originally theorized, based on the lab study, that mixed participants would check permissions at the end of tasks and use the options menu below the album thumbnail to open the permission-modification interface. Instead, what we saw was that participants, regardless of treatment or condition, used the options menu inside an opened album. We anticipate that mixed and under-photo condition participants were clicking on the album and looking at the proximity display. By the time the album opened, they had decided whether there was an error. If they thought a permission error was present, they opened the permission-modification interface; if not, they went on to other actions.

7.3 Proximity-display designs

We tested five different proximity-display designs in the lab and online studies. We apply our observations from the prior sections to better understand how participants interacted

with each style of interface. This section is intended to summarize our findings of each proximity-display design.

7.3.1 Under photo

The under-photo condition kept permission information in close spatial proximity to the users' main focus (the photos) but also used the most screen real estate. It is therefore unsurprising that participants viewing this condition corrected statistically significantly more permissions than they did in the control condition. However, we observed in the lab study that seeing the proximity display in so many places caused some participants to start ignoring the display to the point where they could not remember seeing it. We are concerned that if this design were to become commonly used, people might become habituated to ignoring it. We hypothesize that habituation is one of the reasons that the participants in the under-photo treatment showed a smaller difference between experimental and control conditions than participants in the mixed treatment.

7.3.2 Sidebar


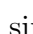

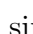
The sidebar condition was not statistically significantly different from any other condition in the number of policy errors corrected or the number of permissions recalled. However, in both the eye-tracker study and the online study the sidebar condition was near significance in the number of permissions corrected. The online-study participants corrected 0.22 more permission errors on average (out of 4) when in the sidebar condition than when in the control condition. If we consider both eye-tracking data (Figure 7.1) and observed behavior (Figure 7.5) we see that, unlike the under-photo and mixed conditions, the sidebar condition does not impact when the participants check or correct permission errors. This condition makes the permission information easier to find, but does not place the display directly in the users' visual path. Consequently, it only helps participants who are looking for permission information. Its advantage over control is that the permissions are easier to find and checking them takes less time and effort. If participants are not looking for the information, or viewing the sidebar for some other reason, they are unlikely to encounter the information. Our results indicate that users in this condition are less likely to identify errors than in other experimental conditions, but it does appear to give more assistance than the control condition and may be worthy of future examination.

7.3.3 Mixed

The mixed interface was designed to combine the best parts of the under-photo and sidebar conditions, based on how users interacted with those conditions. Permission information was placed under album thumbnails so it could be noticed as users entered or closed the album. Under-photo participants in the lab study were observed to primarily check under the album thumbnails, as opposed to under the photo thumbnails. When the album was opened, the proximity display was moved to the sidebar, where it would not interfere with the participants' primary activities. In the lab study we observed that even when

participants identified a permission error midway through the task, they would frequently wait until the end to correct it, potentially forgetting about the error in the process. Eye-tracker study participants in the sidebar condition also checked permissions at the end of the task, a behavior that is enabled by keeping permission information on the sidebar when the album is open. Figure 7.5 suggests that this approach worked well, in that participants who saw the displays tended to open the permission-modification interface early in the task. The shift in when the permission-modification interface was opened may have been because they saw the permission information under the album thumbnail as they opened the album.

7.3.4 Facebook

The Facebook treatment showed no statistically significant difference in the number of permissions corrected between conditions in the online study and was not tested in the other studies. Even the average number of permissions corrected for the two conditions was identical. The lack of difference was likely caused by participants failing to notice or comprehend the proximity display icons. Due to the nature of the errors tested, three of the four tasks with errors displayed a  icon, indicating custom permissions. Facebook uses this icon whenever a user allows a set of groups other than Public, Friends, Friends of Friends, or Private to view the album. A single task showed the  icon, which indicates the album is public. The predominance of the  icon may have put the Facebook interface at a disadvantage. However, the task showing the  icon only saw one participant (0.4%) correct the permissions when they saw the Facebook icon as opposed to the tag icon in the control condition. For comparison: the mixed condition, on this task, had six participants (2.4%) correct permissions when they saw permissions on the proximity display. It is likely that participants simply did not understand the meaning of the icons or did not notice them at all. Also notable is that 89% of participants in the Facebook treatment had previously used Facebook to share photos and were thus previously exposed to the icons.

7.3.5 Audit

We discussed displaying audit information with focus-group participants, piloted several display designs on participants in the eye-tracker and lab studies, and finally conducted a full evaluation in the online study. Our final evaluation showed no statistically significant difference in the number of permissions corrected between control and experimental (audit) conditions. With only 0.007 more permission errors being corrected when participants saw audit information on the display, we are confident that audit displays are not assisting users in identifying errors in this study. Focus-group participants voiced concern that audit data would be unhelpful to them. The lab study's pilot participants were never observed to explicitly check permissions using the audit information on the proximity display. When asked why, participants responded that the information was irrelevant to their goals and not helpful. Audit information is primarily intended to give people feedback about how their implemented policy is being used so they can re-evaluate their ideal and implemented policies and adapt the policies over time. Because this was not our participants' actual

ideal policy, they had no need or interest in re-evaluating, making the data potentially irrelevant to them. We tried to account for this in the online study by including the names of all groups who could view the albums in the displayed information, but this appears not to have helped. However, we feel that the audit proximity display may be more effective in other domains, or when participants' actual data and policies are used. Future researchers and designers should further explore this display.

7.4 Limitations

While we feel that these studies provide valid results that can be applied in other contexts and domains, there are some important limitations the reader should consider.

Role playing

The primary limitation of our study is, we believe, that our participants were challenged to configure policies that were not of their own making and for content that was not their own; this artificiality might have influenced our outcomes.

We chose to use role play with contrived policies because it ensured all participants had a similar experience and that we knew which albums had errors. However, this choice meant that participants were not previously familiar with the ideal policy, and had no real investment in it. It is possible, even likely, that participants might have behaved differently if given an opportunity to work with their own albums.

Perceived risk could also have been an issue. If participants failed to protect a study album, no real harm came to participants. If they do not protect their own albums, there is the potential for actual harm. It is possible that participants might have taken the tasks more seriously, and corrected more permission errors, if the albums had been their own.

The audit treatment in the online study was intended to assist users in both identifying errors and reassessing their prior policy decisions. We showed in the online study that placing audit information on a proximity display did not help participants find policy errors. However, with artificial policies, participants could not really adjust or change the policy as it suited them, so we do not know the impact this display design would have had on actual policies. Additionally, this condition displays a few names of people who have accessed the album in the past, along with the names of the groups those participants are in. The group names should help participants identify errors, but if the names of the users were familiar, they might serve as additional cues — for example, to remind someone that their mother had access to an album or to help them notice that a former employee was still a member of a work group. However, while participants in our study were informed of the names of their friends, family, and co-workers, they may not have internalized the names sufficiently to use them to identify errors.

It would be interesting to re-evaluate our findings on users' own content and policies and in longer-term studies involving repeated user exposure to permissions and the effects of time on their memory.

Photo sharing

Our study used an open-source photo-sharing website system called Gallery 3 and asked participants to conduct photo-manipulation tasks. We selected Gallery 3 because it is relatively unknown to the general populace and could be easily modified. We used version 3.1, that was released in October of 2010 and was a significant departure from the style and user interface of the 2.x versions. We are confident that few, if any, of our participants had prior experience interacting with the 3.1 version of Gallery 3, guaranteeing that all participants received an equal amount of training and experience. However, this also meant that participants were working in an unfamiliar environment. It is possible that, given more time to become familiar with the Gallery 3 interface, participants might have behaved differently.

Our photo-manipulation tasks were designed to simulate users spending time working with their photo albums. The tasks we chose were selected to be plausible and represent tasks an average user might engage in. However, we made no attempt to accurately replicate a typical online photo-editing experience. Our goal was instead to create a scenario that was sufficiently compelling that participants could easily role play it, and that required participants to have the album page open for a similar amount of time across tasks. We feel that our role-playing scenarios are a reasonable approximation of the mindset of users interacting with their online photo albums. However, it is possible that issues such as the length of time spent on the page, or the exact parts of the interface that drew the user's eye, could impact the results of our studies.

Priming

We found that designing a study to test a secondary task, such as permission management, presents inherent difficulties. Notably, participants had to be made aware of what the ideal policy should be, while at the same time not overly biasing them towards fixing permissions. In our studies, participants were directly informed that permission modification and upkeep were a component of the study. By thus informing participants, we effectively primed them to look for permissions, thereby increasing their likelihood of doing so despite our efforts to present the information in a group with similar, irrelevant, information. We anticipate that if we had not primed participants to look for permission errors, we would have seen a lower number of participants finding and correcting errors. We may also have seen a higher difference in the number of permissions checked between the control and experimental conditions.

7.5 Conclusion

We examined the effect of positioning proximity access-control displays near photo albums on participants' ability to notice and correct errors with their access-control permissions. We asked participants to complete several tasks with permission and non-permission actions. We observed that participants in the under-photo and mixed conditions, where access-control information was located under each album thumbnail and under every photo

(under), or under the album thumbnail and on the sidebar (mixed), performed statistically significantly better at checking and fixing errors in albums associated with tasks. We also observed that participants in all conditions tended to change permissions at the beginning and end of tasks, and that some participants were inclined to check all the permissions at once in a single pass. Finally, we observed a high variance between users. Some participants were very inclined to check and correct permissions while others simply forgot about them.

We believe our studies have implications for website-interface design for sites where participants' permission preferences are likely to change over time. It is already the case that empowering end users to effectively manage the privacy of the content they put online is a major issue. Social-networking sites such as Google+ emphasize access control as a way of differentiating themselves from competitors. Our study provides guidance to such sites as to effective means of keeping users more in tune with their policies.

Chapter 8

Designing an access-control study where security is a secondary task

In this chapter we discuss the methodological issues we encountered in our four evaluation studies and how we adjusted our methodologies to emulate security as a secondary task. Simultaneously testing security as a secondary task and clearly conveying to participant their ideal policy is a challenge. In each evaluation study, we tried a slightly different methodology that resulted in differing behavior by participants. This chapter details the different methodological issues we encountered and the solutions that we found to overcome them. Because the chronology of the studies is important for this discussion, in this chapter we refer to the studies by the order in which they were conducted (study 1 through study 4). To assist the reader in understanding the mapping between the studies' names and their order we provide Table 8.1 that shows the order, name, and methodology details of each evaluation study ¹.

Order	Name	Location	Type	Length	Tasks	Conditions
1	Pre-study	Lab	Between-subjects	1 hour	9	3
2	Eye tracker	Lab	Between-subjects	1.5 hours	12	3
3	Lab	Lab	Between-subjects	1.5 hours	16	4
4	Online	Online	Within-subjects	1 hour	16	5

Table 8.1: Order, name, and methodology of each study.

Other researchers have studied security as a secondary task using various approaches [44, 96, 103]. One approach, used by Haake et al. [44], is to conduct a long-term study where participants are made aware that security is a part of the study but the study is run for long enough that users stop focusing on security. Another approach, used by Sunshine et al. [96], is to not make the participants aware of the security nature of the study, but the study design forces participants to engage in a security behavior while trying to complete their primary task. A final approach, used by Wang [103], is to keep participants unaware

¹This chapter is based on a published paper: K. Vaniea, L. Bauer, L. F. Cranor, and M. K. Reiter, *Studying Access Control Usability in the Lab: Lessons Learned From Four Studies*, Proceedings of Workshop on Learning from Authoritative Security Experiment Results, 2012

that the study is about security and give participants the option of whether or not to interact with the security functionality.

To test our hypothesis, that proximity displays help people notice policy errors, we decided to use the last approach. We conducted a lab study where participants performed various photo-management tasks. Depending on the condition, participants were shown proximity displays under the photos and albums, elsewhere on the page, or on a secondary page (control).

When designing the pre-study methodology (study 1), we wanted to meet the following goals: make security a secondary task (Section 8.3), give the participants ownership and responsibility for the albums (Section 8.4), make sure the participants understood the policy they needed to enact (Section 8.5), and develop clear metrics for measuring the outcomes (Section 8.6). Despite careful planning, we encountered methodological issues on every one of these goals.

In this chapter, we discuss the pre-study (study 1), eye-tracker study (study 2), lab study (study 3), and online study (study 4), each of which took into account the methodological issues that arose in the proceeding study. We focus our discussion on aspects of the methodology that tried to accomplish the four goals described above. We describe the difficulties encountered during each study, and changes to the methodology designed to address those difficulties. Through this process, we shed light on the challenges intrinsic to many studies that examine security as a secondary task, and convey a series of lessons that we hope will help other researchers avoid some of the difficulties that we encountered.

8.1 Study goals

The purpose of all four studies was to test the following hypothesis in a photo-sharing system where participants were treating security as a secondary task:

- H1** Users who see information about access-control permission settings on the main interface check and correct permission errors more often than users who have to proactively open a second interface to view permissions.

When designing study 1 to test H1, we wanted to create a study environment that met the following four goals:

Secondary permission task Participants should be in an environment where there is little encouragement to engage in security tasks and the benefits, if any, are not immediate. Users treat security as a secondary task because the benefits of security are often hard to envision, but the cognitive and time costs of engaging in it are immediate [105].

Other researchers who study security technologies have successfully simulated the secondary task mindset in the lab. Whitten and Tygar’s work on email encryption had participants focus on sending and receiving emails while they measured the usability of PGP [107]. Similarly, Sunshine et al. asked participants to find information on websites while studying their reactions to SSL errors [96].

Participant responsibility Participants should feel they are sufficiently responsible for the experimental content to be comfortable making changes they deem necessary. Be-

cause changing permissions is secondary, the framing of the study should make it clear to participants that it is their responsibility to make changes outside the bounds of their primary task.

When replicating the SSL study described above, Sotirakopoulos et al. experienced issues with participants claiming that the lab was a “safe” environment so they behaved differently than they would normally [91]. Whitten and Tygar overcame this issue in their work [107], but doing so requires careful study design.

Ideal policy comprehension Participants should be aware of and comprehend the *ideal policy* – the correct set of permissions for the content. Participants need to have a clear ideal policy associated with the content they are working with. Participants need to be able to consistently decide when the implemented policy is “correct” or “wrong.” If participants are observed to ignore a policy error, we need to have confidence that the error was ignored because participants did not notice the state of the implemented policy rather than because they did not realize the implemented policy did not match the ideal policy.

Effective outcome measurement We need to be able to accurately measure whether participants are noticing and fixing errors. In real world environments, the presence or absence of an error can be very subjective and dependent on context [10, 23, 69]. To accurately test “noticing” errors, we need to be able to differentiate between environments with no errors, environments where participants are not noticing errors, and environments where errors have been noticed.

8.2 General study design

Our initial study design was intended to test the following hypotheses in addition to our main hypothesis H1.

- H2** Participants who see permission information on proximity displays can recall those permissions better than participants who see permission information only if they click to a second page.
- H3** Participants who see proximity displays take no more time, and correct no fewer non-permission errors, than participants who see permission information on a secondary page.
- H4** Users who see permission information under photos and albums notice errors more often than users who see permission information in other spatial locations.
- H5** When a permission is changed to an error state by a 3rd party, users who see permission information under the photos and albums or on the sidebar notice errors more often than users who see permission information only if they click to a second page.
- H6** The type of error, too many permissions or too few, has an effect on the number of errors noticed.

In this chapter we discuss the methodologies of the four evaluation studies briefly. More detailed methodology descriptions are given in Chapter 6. In this section we present the core methodology used in all four studies. In the following sections, we detail the unique

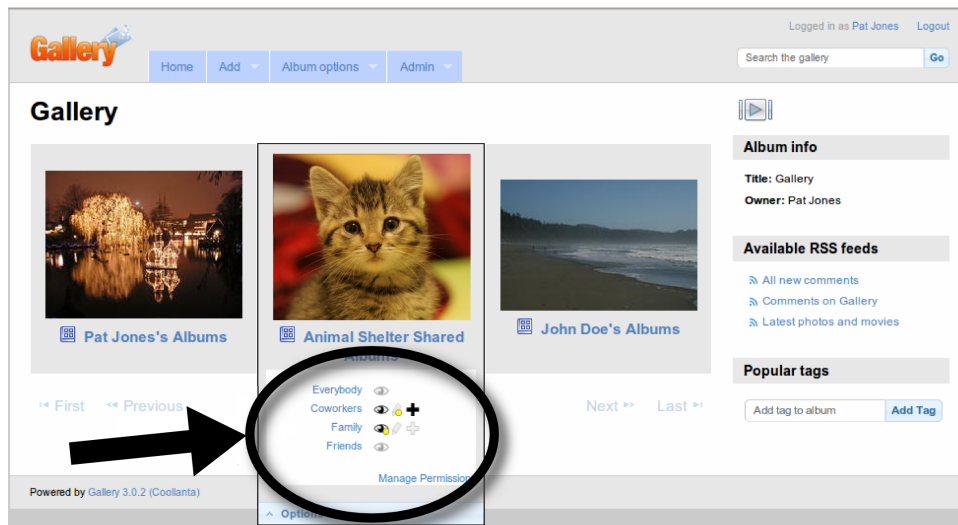


Figure 8.1: Example of proximity display used in studies 1 and 2. The interface for studies 3 and 4 had a slightly different permission display interface design.

methodological choices made in each study to meet the goals described in Section 8.1. We discuss the outcome of the choices and how they informed the methodological choices in the later studies.

The first three studies were between-subjects lab studies and the last was a within-subject online study. All studies used a round-robin assignment to experimental conditions. Participants in all conditions performed the same tasks. Each study had a slightly different set of conditions, but two conditions were present in every study: the control condition was the default interface, which included a link to the interface for changing permissions; the under-photo condition additionally included a proximity display under photo and album thumbnails (Figure 8.1).

Participants were asked to role play [34, 88, 107] the part of Pat Jones, who manages online photo albums using Gallery 3. Role playing is a commonly used method of encouraging user engagement. Whitten et al. successfully use it to encourage participants to view security as a secondary task. Tasks were communicated to participants in email format. In the first three studies, the emails were delivered to participants on paper by the researcher administering the study in the last study, they were shown in an html frame above the website.

Participants started with a training that showed them how to perform several actions on the website including: changing titles, rotating photos, and changing permissions. Participants were asked to perform all actions described in the material to ensure that they understood how to manipulate the interface. In studies 1-3, this training was done on a separate instance of Gallery 3 with fewer albums than the rest of the study. In study 4, the training and the tasks were done on a single Gallery 3 instance.

After the tutorial, participants in study 1 and 2 were given several short warm-up tasks. These tasks were to ensure that participants had understood the training. It also

gave them an opportunity to acclimate to using the interface. Participants in studies 3 and 4 were given 1-2 full task sized warm-up tasks to acclimate to the interface.

The bulk of the studies were composed of a set of tasks presented to participants in sequence. Each task was composed of a set of *errors* – issues with the album that participant were expected to correct to successfully complete a task. A primary error was directly expressed in the email and several additional errors were implied by observable errors such as rotated photos, misspellings, and incorrect permissions. All tasks contained at least one explicit and one implied title, rotate, delete, or organize error intended to distract participants.

Some tasks were *prompted* in that if participants failed to correct any errors, permission related or otherwise, they would be presented with an email pointing out the mistake and asking that it be corrected. *Unprompted* tasks refer either to tasks with no associated prompting or to participant interactions with a task prior to receiving prompting. Participants were unaware of which tasks were prompted until they received a prompt.

Some albums were *changed* midway through the study. First, participants interacted with an album and was made aware of the current state, including the implemented policy. When participant were distracted by another album as part of an unrelated task, the researcher made changes to the initial album. Participant were then instructed to interact with the now changed album again.

Finally, participants filled out a survey that asked them to recall permissions for a selection of albums they worked with, as well as non-task albums with correct and incorrect permissions. For each combination of album, group, and permission participants could answer *True*, *False*, or *Not Sure*. The survey also asked demographic and prior experience questions.

Study 1 was an hour long between-subjects lab study. Participants were given printed training materials that they worked with for about 6 minutes. This was followed by 5 short warm-up tasks that took an average of 8 minutes in total. Participants were then given 8 tasks that took an average of 2.5 minutes each. Tasks appeared in the same fixed order for all participants. Finally, they filled out the survey. There were 5 prompted tasks and 2 changed albums. This study was run on 26 participants and three conditions. It was stopped early because of issues with the methodology.

Study 2 was an 1.5 hours long between-subjects lab study. Participants were given printed training materials that they worked with for about 5.5 minutes. This was followed by 5 short warm-up tasks, that took approximately 8 minutes to complete in total. They were then given 12 tasks to perform, that took an average of 3.5 minutes apiece. Tasks appeared in the same fixed order for all participants. Finally, they were asked to fill out the survey. There were 5 prompted tasks and 3 changed albums. This study was run with 3 conditions and 34 participants; one participant was excluded, that resulted in 11 participants per condition. Further details of this study can be found in Vaniea et al. [102].

Study 3 was a 1.5 hours long between-subjects lab study. Participants were given printed training materials that they worked with for about 5.5 minutes. This was followed by 2 large warm-up tasks taking approximately 13 minutes to complete. They were then given 15 tasks in a random order that took an average of 3.5 minutes apiece. Finally, the survey was verbally administered by the researcher, followed by an unstructured debrief-

ing interview. There were 3 prompted tasks and no changed albums. This study had 2 independent variables: proximity display and permission-modification interface. The proximity display was shown either under the photo (under photo) or not at all (control). The permission-modification interface was either a separate page with all permission settings shown or a dialog with only one album’s permission settings shown. There were 9 pre-study participants and 33 actual participants in this study.

Study 4 was an hour long within-subjects online study conducted on Mechanical Turk. All participants performed training, warm-up, and tasks for both the proximity display condition and the control condition. The order that participants saw the conditions in was assigned round-robin. Participants completed a set of training tasks that took an average of 4 minutes. They completed a warm-up task in an average of 3 minutes. They were then given 7 tasks, with a maximum of 2 minutes to complete each of these tasks. Tasks appeared in the same fixed order for all participants. When finished with both conditions, they were given a survey to fill out that asked questions about both conditions that participants worked with. There was 1 prompted task and 1 changed album per condition. There were 300 pre-study participants and just over 600 actual participants in this study.

8.3 Secondary permission task

Participants should be in an environment where there is minimal encouragement to engage in security tasks, and the benefits, if any, are not immediate.

8.3.1 Study 1

We decided to give participants a primary task that would take the majority of their attention while still being sufficiently open ended enough that they would consider looking for other errors. We communicated the tasks through printed emails because the structure allowed us to give context to the task, such as the ideal policy, without drawing too much attention to it. To prevent users from perceiving permission content as explicit direction, we stated all permission information in passive voice and all primary errors in active voice. For example, the email in Figure 8.2 explicitly asks that the titles be changed, but also implies that the Friends group needs to be able to view the photos. The ideal policy components that could not be expressed passively were embedded in information pages about Pat’s friends, family, and co-workers.

We were concerned about giving participants too much permission *priming* – the amount participants are encouraged to engage in permission behaviors. Every time participants read or interact with permission information, they are being primed to think about permissions. Participants had to be told the ideal policy, which primed them to think about permissions. We compromised by creating three blocks of tasks separated by information pages. Two of the tasks had permission errors, and in the third task permissions were never mentioned. This third task was to give participants time without permission priming.

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: New photos

Yo Pat,

Here are the better photos from the Building Jumping trip last weekend. Could you put them up on your site? Just set it up like any of your other albums. Also could you title the photos with the people in them? I had the red parachute, George had the green one and of course your's was blue.

When you are finished send me back a link so I can forward it to the rest of our friends.

Thanks,
Josh

Figure 8.2: Email from Pat's friend implying that the Friends group needs to be able to view the photos.

To test behavior in the absence of prompting, the first two tasks were unprompted. If participants did not correct permissions on these albums, the researcher did not make them aware of the issue. Participants were first prompted about permissions after the third task. We prompted here to be sure participants knew what the implemented policy for the album was before it was changed by the researcher.

Outcome Participants rapidly deduced that this was an error-finding study and tried to find and correct all the errors. However, none of the participants noticed that the study was solely about permissions. While participants may have been biased to look for errors, only 67% of participants noticed any permission errors without prompting and no participant noticed all the errors. For comparison, 86% of the title errors were corrected.

Over-priming participants to identify and fix errors in general may have caused a control condition behavior we termed "checklisting." Participants who checklisted would reach the end of a task, pause, and appear to go through a mental check list. One participant did this out-loud, listing all the types of errors she had seen in the training material, making sure she had checked all of them before moving on.

Additionally, many participants never obviously consulted the proximity display to determine whether there was an error before opening the permission-modification interface. We hypothesized that since all emails mentioning permissions were associated with albums containing permission errors, participants always needed to open the modification interface and had no need to consult the display.

8.3.2 Study 2

In study 1, all tasks that expressed permission information in the email had permission errors. Effectively there was no “cost” to checking permissions because participants could determine from the email that there was a permission error. To address this concern, we added a new hypothesis:

H7 Participants who see permission information on the main screen are, in the absence of an error, less likely to open the permission-modification screen than users who have to proactively open a second interface to view permissions.

New Read-permission tasks We added three new tasks where the email expressed the ideal policy but the implemented and ideal policies matched, so there was no permission error. After this change, 50% of tasks expressed the ideal policy and had permission errors, 25% of tasks expressed the ideal policy but had no error, and 25% of tasks did not express an ideal policy. Two of the new tasks were prompted. If participants did not obviously check the permissions, the researcher prompted them with an emailed question about the permissions. The new tasks were also intended to test whether participants used the displays to determine the lack of an error (H7).

Outcome The addition of the new tasks appears to have reduced permission priming. We observed no participant engage in checklisting type behavior. Additionally, 53% of participants corrected permissions on 3 or fewer of the 12 tasks before being prompted and no participant corrected all permission errors. In comparison, over 90% of spelling errors were corrected. This suggests that participants were not overly primed to look for permission errors.

The reduction in priming allowed us to observe more subtle issues with our methodology. Participants’ permission-checking frequency was impacted by the different tone and wording of the ideal policy in the task emails. Emails with stronger wording resulted in permissions being checked more frequently by participants in all conditions, and emails with weaker wording were checked less. This meant that while we had a valid study-wide result, we could not compare the permission-identification behavior between tasks. The wording strength added a confounding factor.

8.3.3 Study 3

Reducing the number of tasks with permission errors to 50% and providing ideal-policy information in the absence of errors appeared to cause less checklisting behavior. However, the wording of tasks caused participants to check permissions on some tasks more than others, suggesting that participants did not have consistent priming. In study 3 we wanted the tasks to provide a consistent level of permission priming independent of the presence of a permission error. We also wanted to maintain the “cost” of checking permissions at a 50% chance of there being no error.

One ideal policy We used a single ideal policy that applied to all albums because it 1) better mimicked normal usage where a single user has a consistent set of requirements, 2) was clearer for participants to understand than getting a new policy with every email, and

3) eliminated wording variability since participants would only see one policy. To counter differences in participants' memory, participants were allowed to look back through any piece of paper the researcher gave them, including the page with the policy.

The ideal policy we ultimately selected had five rules, three of which involved access-control permissions. We were concerned that having a single policy that clearly mentioned permissions would overly bias participants to look for permission errors, so we tried the protocol with seven test participants. We found that despite the priming, participants infrequently checked for permission errors but frequently checked for the other types of errors mentioned in the rules.

Consistent task structure Previously, the emails were two paragraphs and important information appeared wherever it was most natural based on the email content. For this study the first paragraph was contextual only, indicating how it related to Pat but containing no vital data. The second paragraph clearly explained the primary error participants were instructed to correct.

Unlike studies 1 and 2, the warm-up tasks in study 3 used the same structure and wording style as the other tasks. Based on observations in the prior studies, the tutorial was sufficient for understanding the system and the warm-up tasks were only necessary for participants to acclimatise to the system and how tasks were presented.

Randomized tasks We decided, with the exception of the warm-up tasks, to randomize the order that tasks were presented in, which tasks had permission errors, and the order the permission errors appeared in. The primary goal was to remove email wording effects, by randomizing if a task has an error as well as randomizing the type of error, we could limit the effect an particular email wording had on permission checking behavior.

Outcome The use of a single ideal policy allowed us to reduce the number of times we presented participants with permission information. Only 11 of the 31 participants checked permissions on more than 50% of the tasks, suggesting that permissions remained a secondary task for the majority of participants.

Our primary concern was that having explicit permission rules expressed in the beginning of the study would overly prime participants to check permissions regularly. The behavior of practice participants suggested that this would not be the case. However, the results of the full study showed that over priming did impact participants.

Our changes to study 2 appeared to eliminate the checklisting behavior observed in study 1 participants, but the design of study 3 brought it back. A graph of the number of tasks where control participants checked permissions shows a non-normal distribution with peaks at 0 and 100. The other conditions showed similar distributions. This suggests that the permission priming affected some participants more than others.

8.3.4 Study 4

In study 3, we saw no difference between conditions because participants corrected all or none of the permissions with few participants in the middle. Using a single ideal policy worked well in study 3, as did the mix of 50% of tasks having permission errors. Because study 4 was within-subject, we decided to use a fixed permission order for easier comparison.

Time limitation We hypothesized in study 3 that providing participants with clearer instructions made it easier for them to know what to do, but the only cost to participants for checking permissions was the time required to perform the check. In real life, that time would be an opportunity trade off since users could be doing something else with that time. In study 4 we decided to limit participants to a maximum of 2 minutes per task, forcing them to value their time and make trade offs. The primary researcher, as an expert user who knew where all the errors were, required a minimum of 1.5 minutes complete each task, so we tried 2 and 3 minute limits on practice participants. We determined that a limit of 2 minutes created the largest differentiation amongst users.

Compensation variation For our practice participants, we were concerned that Mechanical Turk users would not take the tasks seriously and do the minimum to advance through the study. So we offered a bonus based on performance. However, study feedback suggested that participants were deeply concerned that failure to get everything correct meant they would not be paid. They also felt a level of personal responsibility to correct all the errors. So we adjusted compensation to a single rate and explicitly stated that all participants who got more than 25% of the task components correct would be compensated.

Outcome The combination of time limitations and reduction of emphasis on accuracy worked well. Permissions were changed unprompted by 66% of participants. In the under-photo condition, only 4 of the 62 participants corrected all permissions. We also saw a reduction in feedback about the number of tasks participants had correctly completed.

8.4 Participant responsibility

The framing of the study should make participants feel sufficiently responsible for the experimental content to be comfortable making changes they deem necessary, even if the changes are outside the bounds of the errors expressed in the emails.

8.4.1 Study 1

By having participants role play, we were able to inform them that they had a responsibility for some albums by telling them it was part of their job or that their mother regularly relied on them for assistance. We wanted participants to be aware of what types of errors (rotations, spelling, ect.) were within the bounds of the study without overly priming them towards permissions. The tutorial that covered several functionalities of Gallery 3 included permissions. It was followed by five prompted warm-up tasks, two of which involved permissions.

Outcome The open-ended nature of the tasks combined with the imparted responsibility made participants uncertain about how to react to tasks and prompts. For example, after a prompt from Pat’s mother, in which the mother is panicking about seeing a photo of Pat sky diving, one participant simply responded “Sorry Mom.” Another participant asked how old Pat was, then slapped the paper down on the table and declared loudly “I am NOT answering this!”

Some participants did not feel it was their place to change permissions. A couple of participants noticed an error and verbally decided not to correct it because the album belonged to someone else and they expected that the album owner knew what they were doing, even if the permission was odd. Participants were not instructed to talk aloud during the study so we had no way of knowing how many participants noticed an error and chose not to correct it.

8.4.2 Study 2

Based on observations of participants, we theorized that the general uncertainty was caused by a lack of clarity in the task descriptions.

Clearer instructions When observing participants complete the study 1 methodology, we noticed numerous small confusion points that together made participants uncertain about what to do in the study. For example, a warm-up task tells participants that a photo of a poster has an incorrect title but does not specify what the correct title should be. Participants needed to read the title from the photo, but participants became confused. In study 2 we clarified that the titles can be read from the posters in the photos. Another example is from study 1’s task 13, where Pat’s sister apologizes for messing up Mom’s photos and asks Pat to put the photos “back the way you had them.” Participants are supposed to undo changes made by the sister so that the album looks like it did at the end of task 11. Some participants tried to change the album back to what it looked like when they first saw it at the beginning of task 11. We clarified the explanation. When running these tasks on practice participants, we specifically asked them whether these points were clear.

Outcome Participants appeared to have taken responsibility for the albums and considered permissions to be in the bounds of the study. We did not observe any participant choosing to not change permissions due to concern about who owned an album. The clarification in wording resulted in less participant uncertainty over how to handle situations.

8.4.3 Study 3

Directly telling participants that they were responsible for the albums, combined with clear wording, appeared to have caused study 2 participants to sufficiently take responsibility for the albums. In study 3 we tried to keep these themes.

Prompts We initially decided to make only warm-up tasks 1 and 2 prompted tasks to make sure that participants were capable of performing all the actions necessary for the study. As part of the prompting emails, participants are directly told that it is their responsibility to find and fix errors.

After running the protocol on several practice participants, we discovered that around the 5th task, participants would start to become lazy and stop taking responsibility for correcting all the errors. We solved the problem by making task 5 a prompted task. Similar to warm-up tasks 1 and 2, participants were told in the email that fixing errors was their responsibility.

Outcome Participants took responsibility for the albums and considered permissions to be in the bounds of the study. When asked after the study whether they felt they could change permissions, all participants asserted that they felt they were allowed to do so.

Making task 5 a prompted task was very effective in reinforcing participant responsibility. Throughout the study participants would get lazy or careless around this task, receive a strongly worded email from their boss, and immediately start paying more attention. In the debriefing interview, we asked participants about their reaction to this email. Participants said that they realized that the boss would be checking their work so they needed to do a good job.

8.4.4 Study 4

The methodology for study 3 worked well, so we made only minor alterations for study 4. We reduced the strength of wording in the prompted warm-up task so that it simply pointed out the error. Because participants only had eight tasks per condition and were limited to 2 minutes, we decided to not prompt midway through.

Outcome Because study 4 was an online study, we have limited feedback on participants' feeling of responsibility. Participants who gave study feedback expressed a strong desire to get all the tasks correct. The number of permissions and non-permission errors corrected also indicated that participants took responsibility for the albums.

8.5 Ideal policy comprehension

Participants should know the ideal policy associated with the content they are working with.

8.5.1 Study 1

We considered conducting the experiment using participants' own albums and policies but ultimately decided against it. Prior work has shown that participants' ideal policies change over time [69], in reaction to new technology [10], and based on context [23]. Mazurek et al. asked participants to provide ideal policies twice: all at once in a single sitting and by answering the same questions in small batches over the course of a week [69]. They found that the same participants responded with different ideal policies depending upon when asked. We were concerned that participating in our experiment would impact participants' answers concerning their ideal policy, negatively impacting our ability to get an accurate ground truth. Instead we decided to create a fictional ideal policy that would be consistent across all participants.

To make the ideal policy appear less like explicit instructions, we expressed it using passive voice in the emails. However, not all of the ideal policy, particularly who should not see the albums, could be easily expressed in passive voice so some information was presented in instruction pages that described the people participants were about to interact with. To make this information simple to internalize, we created characters. For example: Pat's

mother was described as panicking easily, while Pat was described as enjoying dangerous activities. The instruction sheet commented that Pat generally avoided telling his/her mother about the dangerous activities.

We decided to have two permission warm-up tasks to be certain that participants could accurately both read permissions as well as change the permissions. If they were unable to do so, the researcher provided guidance. The first permission warm-up task simply asked participants whether a particular album was visible to everybody on the Internet or not. The second permission warm-up task asked participants to change the permissions on a specific album.

Outcome Participants seemed to understand the ideal policy without difficulty and participants who made changes tended to make the correct ones. However, we have no way to determine why participants who did not change permissions chose not to do so.

The warm-up task that asked participants to read a permission resulted in participants guessing instead of reading the permission. In the warm-up task, Pat's boss asks whether people at other companies can see a particular album. Participants tended to correctly guess that the album was publicly visible and answered the question without even looking at the screen. We had prepared prompting emails in the event of an inaccurate guess, but had not anticipated that the majority of participants would guess accurately. For the non-control conditions, there was no way to be certain they had guessed since we could not verify if they had looked at the display.

8.5.2 Study 2

Participants seemed to understand the ideal policy in study 1, so we made minimal changes to the way it was presented.

Changed permission-read warm-up task In study 1, participants were guessing that anyone on the Internet could view the album in the permission reading warm-up task. In study 2, we changed the task so that the correct answer was that anyone on the Internet could *not* view the album, thereby making it the opposite of the common guess.

Think-aloud protocol For reasons discussed in following sections, we made study 2 a think-aloud study. A side effect of this decision was that participants had to read all instruction materials and emails out loud, ensuring that all materials, particularly the ideal policy, were read. We were also able to determine when instructions were confusing.

Outcome In warm-up task 2 (read permission) we observed more participants consulting the display to determine what the permissions were instead of opening the permission-modification interface. Participants were still inclined to guess that the album was public but the guesses were now wrong and the researcher was able to prompt them, so every participant understood how to read permissions.

Using a think-aloud protocol forced participants to read all text aloud, thereby ensuring that all materials, including information about the ideal policy, was not skimmed over. Based on the think-aloud statements, participants appear to have understood the ideal policy. However, the protocol had no explicit outcome variable with which to test ideal-policy comprehension.

8.5.3 Study 3

In this study, we decided to present one ideal policy to participants at the beginning instead of presenting the policy in pieces. This was done to provide consistent permission priming (Section 8.3.3). It was also done to promote participant understanding of the ideal policy and make it easier to test that understanding.

Testing ideal policy comprehension Participants in studies 1 and 2 appear to have understood the ideal policy, but we did not measure their comprehension. Study 3 had a single ideal policy so we were able to perform a pre- and post-test of participants' ideal-policy comprehension. The pre-test was administered after the warm-up tasks; participants were asked by a co-worker whether a provided photo was appropriate for the website and if they should do anything when posting it. The post-test was part of the final survey; participants were asked what the permissions for several albums should have been.

Outcome Ideal-policy comprehension was provably high in this study. Participants had no problem remembering the ideal policy and were able to apply it to different situations and albums with high accuracy.

In the pre-test, 78% of participants correctly mentioned permissions for both comprehension questions and only one participant never mentioned permissions. Participants behaved similarly on non-permission comprehension questions. This means that participants were able to 1) recognize that permissions might need to be set for these photos, and 2) correctly apply the ideal policy. Across conditions, participants answered an average of 91% and a minimum of 67% of post-study permission comprehension questions correctly. This shows that the methodology design enabled participants to correctly understand, remember, and apply the ideal policy.

8.5.4 Study 4

As mentioned in Section 8.3.4, we were concerned that the explicit listing of ideal-policy rules in a bulleted list was over-priming participants to look for permission errors. With practice participants in study 4, we experimented with several information-page designs. We conveyed the ideal policy in paragraph form with varying levels of wording intensity and compared that with providing the policy in bullet point form. We found that presenting the policy in bullet point form led to the lowest level of variance and the largest difference in permission correction between conditions.

Outcome In study 3 participants could answer "I do not know" to any comprehension question, but it was rare that they did so. In study 4, 50% of participants answered "I do not know" to at least one comprehension question, but only 4% answered all comprehension questions that way. Of the answered questions, 90% were answered correctly. Interestingly, the design of the information page that conveyed the ideal policy had minimal effect on ideal policy recall. Participants who saw the ideal policy in paragraph form correctly answered approximately 87% of comprehension questions, with minimal variance between designs.

8.6 Effective outcome measurement

We needed to differentiate between environments with no errors, environments where participants were not noticing errors, and environments where errors had been noticed.

8.6.1 Study 1

We chose a lab-study design because it offered us the most control over potential variables. We could control the task design, types of errors, and when errors would appear. By using a role-playing scenario we could also control participants' mindsets when approaching problems.

In order to test our primary hypothesis H1, we needed to detect when a permission error was "noticed." We anticipated that participants who noticed an error were very likely to correct it. For this study, we defined "noticed" as "corrected." The number of people correcting a permission error is a strict subset of the number of people noticing errors and we anticipated a large difference in the number of permissions corrected between the conditions. We were willing to accept that we might not detect participants that chose not to correct a noticed error.

When designing memory questions, we were concerned about participant fatigue leading to questions being guessed at or answered with the fastest answer. To counter this, we limited our questions to six albums and only asked about two of the actions. We also required that all memory questions be answered with True, False, or Not Sure. This was designed to make providing answers the same amount of work as guessing.

Outcome Unfortunately, we did not see a statistically significant difference in the number of permissions corrected between conditions. We also observed participants noticing errors and choosing not to correct them, which was not captured by our definition of "noticed." We considered changing our definition, but determining whether participants had checked the permissions was impossible for participants in the non-control conditions who might or might not have looked at a proximity display. Therefore, while it may be the case that H1 is supported if we define "noticed" as "checked permissions," our lack of measurement fidelity prevented us from testing this.

8.6.2 Study 2

In designing the outcome variables for study 2, we focused on being able to notice when participants checked permissions as well as when they corrected permissions.

Think-aloud and eye tracker Our inability to accurately measure when permissions were noticed but not changed was a major issue with the study 1 methodology. To adjust, we made study 2 a think-aloud study. Study 1 was deliberately not a think-aloud study so we could determine whether participants took an equal amount of time to complete tasks (H3). Think-aloud protocols are known for giving inaccurate timing information. In study 2 we felt that accurate timing information was less important than accurately measuring participants' interactions with the displays.

To assist in measuring if and when participants focused on a display, we decided to use an eye tracker. This data was intended to augment, but not replace, the think-aloud data.

Outcome The think-aloud data enabled us to determine when participants *checked permissions* using the following definition. Control participants were judged to have *checked permissions* if they opened the permission-management interface and the permission was visible on the screen. Participants in the other conditions were judged to have *checked permissions* if they (1) opened the permission-management interface; or (2) read the permission aloud; or (3) clearly indicated through mouse behavior that they were reading the permission display; or (4) pointed at the permission display with their hand while clearly reading the screen. This definition allowed us to measure whether participants paid significant attention to a display.

Data from the eye tracker was less helpful than anticipated. To operate, the eye tracker needed participants' faces to remain in a small area. This is possible for short studies, but our study was 1.5 hours. Participants would shift in their chairs or lean on the desk moving them out of range. We considered prompting participants when they moved outside the required area, but decided this would distract participants and alter their behavior. We tried having participants experiment with the eye tracker before the study so they knew where the optimal area was. This helped, but participants still became distracted by the study and started moving outside the optimal area. While incomplete, the eye tracker data did give us a sense of when participants looked at displays.

8.6.3 Study 3

In study 3, we wanted to get more detailed qualitative data about how and why participants checked permissions. Our definition of "permission checking" from study 2 appeared to be working well so we did not modify it.

Permission modification interface In studies 1 and 2, we observed no difference in implemented policy recall between the conditions (H2). We hypothesized that this was due to the full-sized permission-modification interface. Participants who visited the permission-modification interface tended to change more than one permission indicating that, even in the control condition, they were looking at other permissions. To address this confound, we added the permission-modification interface as an independent variable. The permission-modification interface was either a separate page with the full implemented policy shown or a dialog with only one album's implemented policy shown. We added the following hypothesis:

H8 Participants who see a comprehensive policy-modification interface remember permissions better than participants who see a policy-modification interface that displays a single album.

Post-study memory In studies 1 and 2, we asked participants to answer 128 memory questions about 13 albums, 4 groups and 2 actions (view and add) and saw no statistically significant difference between conditions. In this study we wanted more qualitative data to better understand what people remembered. We decided to verbally administer the

memory questions and elicit free responses. We felt free form answers would provide a better sense of what participants recalled. Once all the memory questions had been asked, the researcher prompted participants about anything they had not yet mentioned. For example, some participants only answered the questions in terms of the view action, so the researcher would ask if they recalled the add or edit action for any of the albums.

When we asked memory recall questions of practice participants who had not checked permissions during the study, we found that they felt embarrassed that they did not know the answer. After several recall questions, they started guessing. To discourage guessing we interleaved the memory recall and comprehension questions. This meant that every participant could, at worst, provide an answer for every other question without having to guess. We found that this discouraged guessing and participants seemed more comfortable admitting that they could not recall the permissions for albums they did not check the permissions on.

Post-study debriefing Once all the questions had been completed we conducted a debriefing interview with participants. In the prior studies, participants had occasionally behaved unexpectedly. Initially we thought this was caused by methodology issues, but some behaviors persisted through different methodologies. In this study, we wanted to get participants' perspective on why they engaged in these behaviors. However, many of the behaviors were short (1-2 seconds long) and we were concerned that participants would not remember why they had made a comment an hour ago. We used a contextual interview approach [40] where participants opened the album they were working with, and the researcher explained the context in which the behavior occurred. The researcher then asked participants questions concerning what they were thinking or why they had done something.

Outcome This study design allowed us to accurately measure and test all the outcome variables we were initially looking for. The only issue was an unknown confounding variable that caused some participants to check permissions frequently and other participants to check them rarely.

The use of a single ideal policy allowed us to observe natural participant behavior that was inhibited by the design of prior studies. In prior methodologies, participants were unable to choose when to check permissions because they did not know the ideal policy until they started a task. With one ideal policy, we observed several participants deciding at a single point in the study to check permissions for every album at once. This behavior was facilitated by the full-permission-modification interface. We found that participants who saw the full interface performed better across several measurements and were more likely to correct permissions regardless of whether they saw the proximity display or not.

The combined use of a single ideal policy, randomized task order, and randomized permission-error order allowed us to notice issues with our definition of permission checking. In the control condition, we reliably determine when the permissions were shown. In the non-control conditions, we only determine when permissions were checked based on participant behavior. In study 3, non-control participants were statistically more likely to check permissions when there was an error than when there was no error. There was no statistical difference for the control participants. This suggests that participants were able to glance at the display and determine if there was an error fast enough not to vocalize

it [100]. This is good news for our display, but it implies that we can only detect when participants *explicitly check* permissions rather than being able to detect every time they noticed permissions. The eye tracker allowed us to determine when they fixate on a display, but similarly did not tell us when they actually noticed the permissions.

The use of contextual interview during the debriefing session was very effective at getting participants to remember their reasoning behind specific actions. In cases where participants could not remember, they were still often able to make an educated guess as to why they would have done an action given their behavior up to that point. While a guess is not as good as remembering, participants' guesses as to reasons behind their actions were more accurate than researchers' educated guesses.

8.6.4 Study 4

The prior studies had a small number of participants, and they exhibited a large between-participant variance, making it difficult to detect differences between conditions. In this study, we wanted to increase the number of participants and account for the variance.

Within subjects In study 3, we observed that some participants internalized the need to check permissions while others did not. In the debriefing interview, the participants who internalized considered it “obvious” and those who did not check permissions appeared to have read the ideal policy and then forgot about permissions. To control for the pre-disposition to pay attention to permissions, we decided to make study 4 a within-subjects study where every participant performs the training and tasks on both the control condition and one of the non-control conditions.

Measuring “noticing” Our hypothesis H1 is that participants in some conditions can “notice” permission errors more frequently than participants in other conditions. In studies 2 and 3, we equated noticing permission errors with checking permissions. However, measuring permission checking requires observation of participants not possible in an online study. Additionally, we showed in study 3 that our measurement of permission checking was, at best, a lower bound for the number of times permissions were actually checked by participants. In study 4, we returned to our definition of “notice” from study 1 where we equated correcting permissions with noticing them. This definition provides only a lower bound, but with the larger number of participants and improvements to the methodology we did not anticipate a problem.

Permission-modification interface In study 3, we observed that participants who saw the permission-modification interface in a dialog had a larger difference in performance between conditions than participants who used the full-page permission-modification interface. Since our main hypothesis H1 is concerned with the impact of proximity displays, not permission-modification interfaces, we decided to use the dialog for study 4.

Outcome Using the stricter definition of “noticed” as “corrected” was effective in that we were able to show statistically significant differences between some of the conditions and control (not all conditions were expected to have a difference). We attribute this to both a larger number of participants and clearer, more tested, study materials.

Similar to study 1, we had a limited ability to measure why participants did or did not make changes to permissions. However, we collected extensive logs that we were able

to compare to behaviors observed in prior studies, allowing us to imply what users were doing.

8.7 Discussion

We discussed the methodologies of four studies designed to test our hypothesis. When designing our initial study, we tried to account for anticipated methodology issues. Our initial design succeeded in some aspects and was lacking in others. Subsequent studies were adjusted to account for observed issues.

Secondary-permission task Users treat security as a secondary task because the benefits of security are hard to envision, whereas the costs of engaging in it are immediate [105]. In our studies, we did not want to incentivize participants to check permissions, so we tried to balance the amount of priming with the cost of checking. We successfully managed priming on study 2 and 4, but in studies 1 and 3 we over-primed, first by mentioning permissions too frequently and then by using strong wording to express the ideal policy without forcing participants to consider trade-offs. We increased the immediate cost of checking permissions in studies 2 and 3 by adding tasks where the permissions were already correct and checking them cost unnecessary time and effort. We further increased the cost in study 4 by adding a time limitation that forced participants to make trade-offs. We found that at least 50% of the tasks needed to have no permission error in order to give checking a high cost compared to the benefit.

Participant responsibility Role-playing was very effective in making participants feel responsible for albums that belonged to Pat. Our main issue was when we asked participants to be responsible for albums that belonged to people such as Pat’s mother. We countered this issue in the second study by making it clearer that others trusted Pat to make changes.

Ideal policy comprehension We tried two methods of expressing the ideal policy to participants. The first was to have a different policy for each album. The policy was expressed using passive voice in the emails (studies 1 and 2). The second way was to have a policy that applied to all the albums. The policy was expressed using direct wording at the beginning of the study (study 3 and 4). Both methods sufficiently communicated the policy to participants. The per-album policy gave participants less priming towards fixing permissions but was difficult to make consistent. The study-wide policy over-primed some participants to look for permission errors, but provided consistent priming to all participants on all tasks.

Effective outcome measurement Our primary issue with measuring the study outcome was defining and testing participants’ ability to “notice” permission errors. In the first study we defined “notice” as changing permissions, but this definition was insufficiently precise to measure the difference between conditions. In later studies, we changed our definition of “notice” to checking the permissions for errors. This definition allowed us to observe whether participants were looking for errors independently of whether they found the error or decided to fix it.

In conclusion, we presented the methodologies of four studies and discussed the decisions

and outcomes of each study. We were able to describe our methodological successes and difficulties in terms of our four goals: 1) secondary permission task, 2) participant responsibility, 3) ideal policy comprehension, and 4) effective outcome measurement. Through this process, we have shed light on the challenges intrinsic to many studies that examine security as a secondary task.

Chapter 9

Conclusion

This thesis addresses the issue of helping users better understand their implemented policy and better identify policy errors. We find that people use a variety of methods to control the security of their resources and data. Many of these methods are predicated on users being aware both of their implemented policy, i.e., of what could happen, and of how that policy has been used, i.e., what has happened. To improve users' understanding of their implemented policy and their ability to maintain it over time, we proposed the use of proximity displays — small interface components spatially located near the data elements (or near a representation of data, e.g., file name in a file manager or thumbnail photo in a photo album) that contain information about who has or who could access the data. We applied the concept to a photo-sharing website where users treat security as a secondary task. We then tested the following hypothesis:

Users of a system that includes proximity information displays of access-control information will implement policies that result in grant/deny actions that better match their preferences than will users of a system where access-control information is available only on a secondary interface.

To test the hypothesis we conducted focus group studies to gauge user reactions, and empirical evaluations to test the effectiveness of the different proximity-display designs at improving users' error identification and understanding of their implemented policy. In the focus groups we found that, for the personal photo domain, users liked the idea of making privacy-policy settings appear in close proximity to the photos. However, participants had a strong association between seeing detailed information about who had viewed photos in the past and stalking behavior. The evaluation studies showed that participants who saw proximity displays with comprehensive permission information that could be easily glanced at were better able to identify access-control policy errors. Participants who saw displays that exclusively used icons, were located on the sidebar, or presented information about who had previously viewed the photos, showed no improvement over users who saw permission settings only on a secondary interface. Our studies suggest that using proximity displays to show the implemented policy can significantly help users identify permission errors.

While the proximity displays appear to help people find permission errors, they seem to have no effect on implemented policy recall.

The hypothesis is partially supported by our work. People who see some proximity-display designs are more likely to notice errors than when they have to visit a secondary interface to see the permission information. However, we have been able to find no effect on implemented policy recall. We conclude that proximity displays are a promising approach to assist users in error detection in privacy policies, but have minimal impact on implemented policy recall, at least in the form we investigated.

9.1 Contributions

The original contributions of this thesis provide a guide for both researchers trying to design lab studies in this space and interface designers wanting to create interfaces that increase policy awareness and enable people to decrease policy errors.

Evaluated hypotheses: Our evaluation of proximity displays quantitatively examined three points (Chapter 7):

Correcting/checking permissions: We find that users who see permission information on a proximity display check and correct permission errors more often than users who see permission information on a secondary page.

Showing the display under the album and photo thumbnails, or under the album thumbnail and on the album sidebar, appeared to have the largest effect on participants. They appeared to check the display as they opened the album or after they had closed the album. Conditions that placed displays under the album thumbnail showed a statistically significant improvement in participants' ability to identify errors.

Permission recall: Participants who see permission information on proximity displays do not recall those permissions better than participants who see permission information only if they click to a second page.

Negative effects: Participants who see proximity displays correct no fewer non-permission errors than participants who see permission information on a secondary page.

Understanding user behavior and sentiment: We observed in both our focus group studies (Chapter 4) and our evaluation studies (Chapter 7) that some users are more concerned about permission settings than others. In our focus group studies, some participants strongly felt that their privacy settings did not really matter because websites would likely lose or expose their photos anyway. This difference caused them to view permission information as unimportant.

We observed in the lab and online studies that people in the control condition check permissions primarily at the end of tasks and rarely at the beginning. Participants seeing proximity displays under the album thumbnail tend to check permissions at the beginning of tasks rather than at the end.

Audit information: Showing people detailed information about who has previously seen their personal photos was not well liked by users (Chapter 4). Participants felt that

making people privy to so much detailed information encouraged stalking behavior. Online study participants did not show an improved ability to notice and correct permissions over control when shown audit information (Chapter 7). However, when asked about a document-sharing system, focus-group participants liked the idea of seeing audit information and considered it a useful component of online document management.

Methodology: studying security as a secondary task: Designing a methodology that encourages participants to treat security as a secondary task while at the same time imparting the goal permission state, is challenging (Chapter 8). Our analysis of the methodological issues we encountered, their causes, and how to overcome them, is a valuable tool for future researchers in this domain.

9.2 Future work

Future work falls into three categories: developing more effective proximity display designs, understanding what causes people to look for errors in their access-control policies, and exploring domains beyond personal photo management.

9.2.1 Proximity-display design

In our studies we have explored several designs and spatial placements for proximity displays, but we have only taken an initial look at the space of possible designs.

- **Additional privacy settings:** We limited our analysis to showing permission information related to what other people could do, or had done. However, there are many other privacy settings that could be placed on proximity displays. For example, on Facebook users can control what information is available to their friends when their friends use applications, as opposed to what their friends can see normally.
- **Display designs:** We explored only a small portion of the possible designs for proximity displays. There are many different ways to display privacy information in ways that can be easily glanced at [97]. In particular, we would like to explore the effectiveness of using different styles of icons and other compact policy representations.
- **Error detection at a glance:** In this and other work [101, 108] we see that people glance at information displays and if they do not detect an issue quickly, they assume there is no issue and move on. A proximity display needs to show people enough information that they can accurately identify an error at a glance. If too little information is shown, users may inaccurately decide that there is no error. The question is what data best assists users in identifying errors and how much of it is necessary. The displays we proposed use a non-trivial amount of screen real estate. We would like to know how compact the display can be made before its effectiveness begins to decrease.
- **Display locations:** We showed that placing displays under every photo and album, or under every album thumbnail and on the sidebar, helped participants identify

errors. However, we anticipate that with more participants (power) we might see a significant effect when permissions are displayed just on the sidebar. Additionally, there may be other places and times when showing the proximity display might be more effective.

- **Habituation:** Many of our proximity-display designs were unfamiliar to our participants. As we have seen previously, displays that work initially may stop working when participants become habituated to ignoring them [91, 96]. Proximity displays should be tested for an extended time period to determine whether users continue to interact with them after they have become familiar with the interface.

9.2.2 Understanding policy error-identification behavior

We have shown that participants tend to check permissions at the end of tasks, and that exposing them to proximity displays causes them to check at the beginning of tasks more frequently. We also observed that some people are more inclined to check permissions for errors than other people. Post-study survey answers, and information from the focus groups, suggest that peoples’ assumptions about whether privacy settings on websites will be effective at protecting their content impacts their permission checking behavior. We would like to further explore this observation and determine whether people’s mental models of website behavior really do impact permission checking.

9.2.3 Exploring proximity displays in other domains

In this work we looked at proximity displays in the domain of online photo sharing. However, we believe that proximity displays could be effective in helping end users manage their access-control policies in a variety of domains. The following are examples of possible domains.

Social networking A clear extension to this work is to test proximity displays in a social-networking site context. Social-networking sites, such as Facebook, are creating increasingly more complex privacy policies that users can configure. In future work, we would like to explore how these settings could be incorporated into the proximity-display design.

There is also the issue of users’ understanding of implemented policy: while our studies did not indicate that proximity displays improved participants’ ability to recall permissions, they did make it easier for participants to find, and check, implemented permissions when they were interested. Placing setting information on the proximity display may improve people’s understanding of what settings are available to be manipulated. In addition to not having an accurate understanding of their own permission settings, people are not always aware of all the settings that are available [56]. For example, users may not attempt to opt-out of marketing data being sold if they are not aware that opting out is an option.

Document sharing Managing document sharing in an organization is an issue corporate IT departments are struggling with [101]. Documents are easy to create and share, but people will resort to email and USB drives if the corporate document-sharing interface is too restrictive or unusable. While convenient, these technologies are less secure and are more likely to be lost or compromised than the company’s servers. Additionally, as the internal structure of an organization changes, the access-control policy does not always change with it, leaving people with too much or too little access than is necessary to do their jobs. Proximity displays could help people keep up with the changes by helping them identify permission errors and enabling them to easily determine who can see each document.

Health care The domain of health care is interesting in that emergency personnel need immediate access to health care records for safety reasons, and the data in medical files is generally considered privacy sensitive. To help address this issue *ex-post* access-control management solutions have been explored by researchers. One such approach is termed break-the-glass in reference to breaking glass to access items such as fire extinguishers during a fire. The idea is to give people the access they will need in the majority of situations, and in case of an emergency they can perform a special request that immediately grants them access but is also flagged for audit [20]. In this way, emergency personnel can access any file they feel to be necessary, but know that they will have to be able to justify the access later.

Systems like break-the-glass are potentially effective when oversight is practical. However, some organizations are too small to employ auditors, or the people doing the auditing do not have sufficient understanding of the incidents to properly judge the appropriateness of the accesses. Proximity displays are similar to technologies like break-the-glass in that they help health care professionals maintain security on the files through *ex-post* control. A system using proximity displays could allow anyone access to any file, but also show that access attempt to anyone else interacting with the file. Effectively this would allow access attempts to be judged by peers, not just management.

Chapter 10

Bibliography

- [1] Gallery 3. Accessed on: July 2012, <http://gallery.menalto.com/>. 5.2
- [2] Generation Y online security survey. Technical report, TRU Research, 2010. 4.3.1
- [3] Facebook settles FTC charges that it deceived consumers by failing to keep privacy promises, November 2011. Accessed on: July 2012, <http://ftc.gov/opa/2011/11/privacysettlement.shtm>. 4.3.1
- [4] Anne Adams and Martina Angela Sasse. Users are not the enemy. In *Communications of the ACM*, volume 42, pages 40 – 46, 1999. 2.2.4
- [5] Shane Ahern, Dean Eckles, Nathaniel S. Good, Simon King, Mor Naaman, and Rahul Nair. Over-exposed?: Privacy patterns and considerations in online and mobile photo sharing. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2007. 2.2.1
- [6] Dawn M. Appelli, Akash G. Desai, Andrew P. Moore, Timothy J. Shimeall, Elise A. Weaver, and Bradford J. Willke. Management and education of the risk of insider threat (MERIT): Mitigating the risk of sabotage to employers’ information, systems, or networks. Technical Report CMU/SEI-2006-TN-041, CERT, Software Engineering Institute at Carnegie Mellon University and CyLab, 2007. 2.2.4
- [7] Dixie B. Baker. Fortresses built upon sand. In *Proceedings of the workshop on New security paradigms*, 1996. 2.2.3
- [8] Lujo Bauer, Scott Garriss, and Michael K. Reiter. Distributed proving in access-control systems. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2005. 2.3.1
- [9] Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, and Kami Vaniea. Lessons learned from the deployment of a smartphone-based access-control system. In *Proceedings of the Symposium on Usable Privacy and Security*, 2007. 1, 2.2.6, 2.3, 2.3.1, 2.3.2, 2.4
- [10] Lujo Bauer, Lorrie Faith Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. A user study of policy creation in a flexible access-control system. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*,

2008. 1, 2.2.3, 2.3.1, 2.3.2, 2.4, 2b, 8.1, 8.5.1
- [11] Lujo Bauer, Lorrie Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. Real life challenges in access-control management. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2009. 1, 2.2.6
 - [12] Lujo Bauer, Lorrie F. Cranor, Robert W. Reeder, Michael K. Reiter, and Kami Vaniea. Effects of access-control policy conflict-resolution methods on policy-authoring usability. Technical Report CMU-Cylab-09-006, Cylab, Carnegie Mellon University, March 2009. 2.3.1, 5.4.3
 - [13] Allan Beaufour and Philippe Bonnet. Personal servers as digital keys. In *Proceedings of the IEEE International conference of Pervasive Computing and Communications*, 2004. 2.3.1
 - [14] Victoria Bellotti and Abigail Sellen. Design for privacy in ubiquitous computing environments. In *Proceedings of the conference on European conference on Computer-Supported Cooperative Work*, 1993. 2.2.2
 - [15] Michael Benisch, Patrick Kelley, Norman Sadeh, and Lorrie Faith Cranor. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Journal of Personal and Ubiquitous Computing*, pages 1–16, 2010. 2.2.6
 - [16] Andrew Besmer and Heather Richter Lipford. Moving beyond untagging: photo privacy in a tagged world. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2010. 1, 2.2.1
 - [17] Hugh Beyer and Karen Holtzblatt. *Contextual Design: Defining customer-centered systems*. Morgan Kaufmann Publishers, 1998. 4.3, 6.3.4
 - [18] Bob Blakley. The Emperor’s old armor. In *Proceedings of the workshop on New security paradigms*, 1996. 2.3.3
 - [19] Carolyn A. Brodie, Clare-Marie Karat, and John Karat. An empirical study of natural language parsing of privacy policy rules using the sparcle policy workbench. In *Proceedings of the Symposium on Usable Privacy and Security*, 2006. 2.3.1
 - [20] Achim D. Brucker and Helmut Petritsch. Extending access control models with break-glass. In *Proceedings of the ACM symposium on Access Control Models and Technologies*, 2009. 2.2.6, 9.2.3
 - [21] A.J. Brush and Kori Inkpen. Yours, mine and ours? Sharing and use of technology in domestic environments. In *Proceedings of the international conference on Ubiquitous computing*, 2007. 2.2.3
 - [22] L. Jean Camp, Cathleen McGrath, and Alla Genkina. Security and morality: A tale of user deceit. *Models of Trust for the Web*, 22, 2006. 4.3.1
 - [23] Sunny Consolvo, Ian E. Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. Location disclosure to social relations: Why, when, & what people want to share. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2005. 2.2.2, 8.1, 8.5.1
 - [24] Fergus I.M. Craik and Robert S. Lockhart. Levels of processing: A framework for

- memory research. *Journal of Verbal Learning and Verbal Behavior*, 11:671–684, 1972. 7.2.2
- [25] Lorrie Faith Cranor. Privacy policies and privacy preferences. In L. F. Cranor and S. Garfinkel, editors, *Privacy and Usability*. O’Reilly, 2005. 1, 3.1.2
 - [26] Lorrie Faith Cranor. A framework for reasoning about the human in the loop. In *Proceedings of the conference on Usability, Psychology, and Security*, 2008. 2.5, 2.5.2, 7.2.1
 - [27] Justin Cranshaw, Jonathan Mugan, and Norman Sadeh. User-controllable learning of location privacy policies with gaussian mixture models. In *Proceedings of the AAAI conference on Artificial Intelligence*, 2011. 2.2.6
 - [28] Brinda Dalal, Les Nelson, Diana Smetters, and Nathaniel Good. Ad-hoc guesting: When exceptions are the rule. In *Proceedings of Usability, Psychology, and Security*, 2008. 2.2.4, 2.2.6
 - [29] Rogério de Paula, Xianghua Ding, Paul Dourish, Kari Nies, Ben Pillet, David Redmiles, Jie Ren, Jennifer Rode, and Roberto Silva Filho. Two experiences designing for effective security. In *Proceedings of the Symposium on Usable privacy and security*, 2005. 1, 2.3.1, 2.4
 - [30] Paul Dourish, E. Grinter, Jessica Delgado de la Flor, and Melissa Joseph. Security in the wild: user strategies for managing security as an everyday, practical problem. *Journal of Personal Ubiquitous Computing*, 8:391–401, 2004. ISSN 1617-4909. 1, 2.2.2, 2.2.3, 2.4
 - [31] W. Keith Edwards, Erika Shehan Poole, and Jennifer Stoll. Security automation considered harmful. In *Proceedings of the IEEE New Security Paradigms workshop*, 2007. 2.2.6
 - [32] Serge Egelman. *Trust Me: Design Patterns for Constructing Trustworthy Trust Indicators*. PhD thesis, Carnegie Mellon University, 2009. CMU-ISR-09-110. 1.1, 2.1
 - [33] Serge Egelman, A.J. Brush, and Kori Inkpen. Family accounts: A new paradigm for user accounts within the home environment. In *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 2008. 2.2.3
 - [34] Serge Egelman, Andrew Oates, and Shriram Krishnamurthi. Oops, i did it again: Mitigating repeated access control errors on facebook. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2011. 8.2
 - [35] Willis D. Ellis. *A Source Book of Gestalt Psychology*. Psychology Press, 1999. 3.3, 5.3.1
 - [36] David Ferraiolo, Janet A. Cugini, and D. Richard Kuhn. Role-based access control (rbac): Features and motivations. In *Proceedings of Annual Computer Security Application conference*, 1995. 2.2.4, 2.3
 - [37] David F. Ferraiolo, Dennis M. Gilbert, and Nickilyn Lynch. An examination of federal and commercial access control policy needs. In *National Computer Security conference*, 1993. 2.2.4

- [38] Shirley Gaw, Edward W. Felten, and Patricia Fernandez-Kelly. Secrecy, flagging, and paranoia: adoption criteria in encrypted email. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2006. 2.2.5, 4.3.1
- [39] Virgil Gligor. Characteristics of role-based access control. In *Proceedings of the ACM workshop on Role-based access control*, 1996. 2.3
- [40] D Godden and A.D Baddeley. Context-dependent memory in two natural experiments: on land and under water. *British Journal of Psychology*, 66:325–331, 1975. 6.3.2, 8.6.3
- [41] Grant Gross. Former gov’t worker sentenced for passport snooping, March 2009. Accessed on: August 5th 2009, <http://www.networkworld.com/news/2009/032309-former-govt-worker-sentenced-for.html>. 2.2.4
- [42] Joshua B. Gross and Mary Beth Rosson. Looking for trouble: understanding end-user security management. In *Proceedings of the Symposium on Computer Human Interaction for the Management of Information Technology*, 2007. 2.2.4
- [43] Alejandro Gutierrez, Apeksha Godiyal, Matt Stockton, Michael LeMay, Carl A. Gunter, and Roy H. Campbell. Sh@re: negotiated audit in social networks. In *Proceedings of the IEEE International conference on Systems, Man and Cybernetics*, 2009. 2.3.3
- [44] Joerg M. Haake, Anja Haake, Till Schümmer, Mohamed Bourimi, and Britta Landgraf. End-user controlled group formation and access rights management in a shared workspace system. In *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 2004. 8
- [45] Cormac Herley. So long, and no thanks for the externalities: The rational rejection of security advice by users. In *Proceedings of New Security Paradigms workshop*, 2009. 2.2.4
- [46] Airlie House and Virginia Warrenton. CRA conference on grand research challenges in information security & assurance, November 2003. Accessed on: July 2012, <http://www.cra.org/Activities/grand.challenges/security/home.html>. 2
- [47] Trent Jaeger, Antony Edwards, and Xiaolan Zhang. Managing access control policies using access control spaces. In *Proceedings of the ACM symposium on Access control models and technologies*, 2002. 2.3, 2.3.3
- [48] Lukasz Jedrzejczyk, Blaine A. Price, Arosha K. Bandara, and Bashar Nuseibeh. On the impact of real-time feedback on users’ behaviour in mobile location-sharing applications. In *Proceedings of the Symposium on Usable Privacy and Security*, 2010. 2.2.2, 2.3.2
- [49] Xin Jin, Ram Krishnan, and Ravi Sandhu. A unified attribute-based access control model covering dac, mac and rbac. In *Proceedings of the Annual WG 11.3 conference on Data and Applications SEcurity and Privacy*, 2012. 2.2.6
- [50] Maritza L. Johnson, Steven M. Bellovin, Robert W. Reeder, and Stuart E. Schechter. Laissez-faire file sharing: Access control designed for individuals at the endpoints.

- In *Proceedings of the workshop on New security paradigms workshop*, 2009. 2.3.1
- [51] Sara Kaemer and Pascale Carayon. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. In *Applied Ergonomics*, 2007. 2.2.4
 - [52] Apu Kapadia, Geetanjali Sampemane, and Roy H. Campbell. Know why your access was denied: regulating feedback for usable security. In *Proceedings of the ACM conference on Computer and Communications Security*, 2004. 2.4
 - [53] Clare-Marie Karat, John Karat, Carolyn Brodie, and Jinjuan Feng. Evaluating interfaces for privacy policy rule authoring. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2006. 2.3.1
 - [54] Patrick Gage Kelley, Lucian Cescă, Joanna Bresee, and Lorrie Faith Cranor. Standardizing privacy notices: an online study of the nutrition label approach. In *Proceedings of the 28th international conference on Human factors in computing systems, CHI '10*, 2010. 3.1.2
 - [55] Timothy Kelly, Suzanne Lien, L. Jean Camp, and Douglas Stebila. Self-identified experts lost on the interwebs. In *Proceedings of the Learning from Authoritative Security Experiment Results (to appear)*, 2012. 2.1
 - [56] Jennifer King, Airi Lampinen, and Alex Smolen. Privacy: Is there an app for that? In *Proceedings of the Symposium on Usable Privacy and Security*, 2011. 2.2.1, 9.2.3
 - [57] Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujio Bauer, Lorrie Faith Cranor, Nitin Gupta, and Michael Reiter. Tag, you can see it!: using tags for access control in photo sharing. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2012. 1, 2.2.2, 2.2.6
 - [58] Brian Krebs. Court rules against teacher in myspace 'drunken pirate' case, December 2008. Accessed on: July 2012, http://voices.washingtonpost.com/securityfix/2008/12/court_rules_against_teacher_in.html. 1
 - [59] Ponnurangam Kumaraguru, Justin Cranshaw, Alessandro Acquisti, Lorrie Faith Cranor, Jason Hong, Marry Ann Blair, and Theodore Pham. School of phish: A real-world evaluation of anti-phishing training. In *Proceedings of the Symposium on Usable Privacy and Security*, 2009. 3.1.2
 - [60] B.W. Lampson. Protection. *ACM SIGOPS Operating Systems Review*, 8(1):18–24, 1974. 2.3
 - [61] Barry Leibowitz. Facebook blunder invites 15,000 to teen's 16th birthday party; 100 cops show up, too, June 2011. Accessed on: July 2012, http://www.cbsnews.com/8301-504083_162-20069457-504083.html. 1
 - [62] Eric Lieberman and Robert C Miller. Facemail: showing faces of recipients to prevent misdirected email. In *Proceedings of the Symposium on Usable Privacy and Security*, 2007. 2.1
 - [63] Linda Little, Elizabeth Sillence, and Pam Briggs. Ubiquitous systems and the family:

- thoughts about the networked home. In *Proceedings of the Symposium on Usable Privacy and Security*, 2009. 2.2.3
- [64] Mary Madden and Aaron Smith. Reputation management and social media. Technical report, Pew Internet & American Life Project, 2010. 2.2.2
 - [65] Michelle Madejski, Maritza Johnson, and Steven M. Bellovin. The failure of on-line social network privacy settings. Technical Report CUCS-010-11, Department of Computer Science, Columbia University, 2011. 1, 2.2.1, 2b
 - [66] Roy A. Maxion and Robert W. Reeder. Improving user-interface dependability through mitigation of human error. *Int. J. Hum.-Comput. Stud.*, 63:25–50, 2005. 2.3.1, 3.1.2, 5.4.3
 - [67] Alain Mayer, Avishai Wool, and Elisha Ziskind. Fang: A firewall analysis engine. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2000. 2.3
 - [68] Michelle L. Mazurek, J.P. Arsenault, Joanna Breese, Nitin Gupta, Iulia Ion, Christina Johns, Daniel Lee, Yuan Liang, Jenny Olsen, Brandon Salmon, Richard Shay, Kami Vaniea, Lujo Bauer, Lorrie Faith Cranor, Gregory R. Ganger, and Michael K. Reiter. Access control for home data sharing: Attitudes, needs and practices. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2010. 2.2.2, 2.2.3, 2.4
 - [69] Michelle L. Mazurek, Peter F. Klemperer, Richard Shay, Hassan Takabi, Lujo Bauer, and Lorrie Faith Cranor. Exploring reactive access control. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2011. 1, 2.2.2, 2.3.2, 8.1, 8.5.1
 - [70] Sylvie Noël and Jean-Marc Robert. Empirical study on collaborative writing: What do co-authors do, use, and like? *Comput. Supported Coop. Work*, 13:63–89, 2004. 2.2.4
 - [71] D.A. Norman. *The design of everyday things*. Basic Books New York, 2002. 2.3.2
 - [72] National Academy of Engineering. Grand challenges for engineering: Secure cyberspace. Accessed on: July 2012, <http://www.engineeringchallenges.org/cms/8996/9042.aspx>. 2
 - [73] Antti Oulasvirta. Finding meaningful uses for context-aware technologies: The humanistic research strategy. In *Proceedings of Computer Human Interaction*, 2004. 2.2.2
 - [74] Sameer Patil and Jennifer Lai. Who gets to know what when: configuring privacy permissions in an awareness application. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2005. 2.2.2
 - [75] Erika Shehan Poole, Marshini Chetty, Tom Morgan, Rebecca E. Grinter, and W. Keith Edwards. Computer help at home: methods and motivations for informal technical support. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2009. 2.2.3
 - [76] Dean Povey. Optimistic security: A new access control paradigm. In *Proceedings of*

- New Security Paradigms workshop*, 1999. 2.2.6, 2.3, 2.3.3
- [77] E. Rader, R. Wash, and B. Brooks. Stories as informal lessons about security. In *Proceedings of the Symposium on Usable Privacy and Security*, 2012. 2.2.2, 2.2.3
 - [78] Robert W. Reeder, Clare-Marie Karat, John Karat, and Carolyn Brodie. Usability challenges in security and privacy policy-authoring interfaces. In *Human-Computer Interaction*, 2007. 2.3.1
 - [79] Robert W. Reeder, Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, Kelli Bacon, Keisha How, and Heather Strong. Expandable grids for visualizing and authoring computer security policies. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2008. 1
 - [80] Robert W. Reeder, Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, Kelli Bacon, Keisha How, and Heather Strong. Expandable grids for visualizing and authoring computer security policies. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2008. 2.3.1, 3.1.2, 4.3.4, 5.3.1, 5.4.1
 - [81] Robert W. Reeder, Lujo Bauer, Lorrie Faith Cranor, Michael K. Reiter, and Kami Vaniea. More than skin deep: Measuring effects of the underlying model on access-control system usability. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2011. 5.3.1
 - [82] E. Rissanen, B. Firozabadi, and M. Sergot. Towards a mechanism for discretionary overriding of access control. In *Security Protocols*, 2002. 2.2.6, 2.3, 2.3.3
 - [83] Norman Sadeh, Jason Hong, Lorrie Faith Cranor, Ian Fette, Patrick Gage Kelley, Maduh Prabaker, and Jinghai Rao. Understanding and capturing people’s privacy policies in a mobile social networking application. *Journal of Personal and Ubiquitous Computing*, 13(6), 2009. 2.2.6
 - [84] Brandon Salmon, Frank Hady, and Jay Melican. Learning to share: A study of sharing among home storage devices. Technical Report CMU-PDL-07-107, Carnegie Mellon University Parallel Data Lab, October 2007. 2.2.3
 - [85] Brandon Salmon, Steven W. Schlosser, Lorrie Faith Cranor, and Gregory R. Ganger. Perspective semantic data management for the home. In *Proceedings of the USENIX conference on File Storage Technologies*, 2009. 2.2.2, 2.2.3
 - [86] Jerome H. Saltzer and Michael D. Schroeder. The protection of information in computer systems. *IEEE, Proceedings*, 63:1278–1308, 1975. 2.3
 - [87] Bruce Schneier and Marcus Ranum. Schneier-ranum face-off: Is perfect access-control possible?, September 2009. Accessed on: July 2012, http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1365957,00.html. 2.2.4
 - [88] Steve Sheng, Mandy Holbrook, Ponnurangam Kumaraguru, Lorrie Faith Cranor, and Julie Downs. Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the annual SIGCHI conference on Human factors in computing systems*, 2010. 8.2

- [89] Sara Sinclair, Sean W. Smith, Stephanie Trudeau, M. Eric Johnson, and Anthony Portera. Information risk in the professional services - field study results from financial institutions and a roadmap for research. Technical report, Dartmouth College, 2007. 2.2.5
- [90] D. K. Smetters and Nathan Good. How users use access control. In *Proceedings of the Symposium on Usable Privacy and Security*, 2009. 2.2.5
- [91] Andreas Sotirakopoulos, Kirstie Hawkey, and Konstantin Beznosov. On the challenges in usable security lab studies: Lessons learned from replicating a study on ssl warnings. In *Proceedings of the Symposium on Usable Privacy and Security*, 2011. 2.1, 3.1.3, 8.1, 9.2.1
- [92] Hanna Stelmazewska, Bob Fields, and Ann Blandford. The roles of time, place, value and relationships in collocated photo sharing with camera phones. In *Proceedings of the British HCI Group Annual conference on People and Computers: Culture, Creativity, Interaction - Volume 1*, 2008. 2.2.2
- [93] Gunnar Stevens and Volker Wulf. A new dimension in access control: Studying maintenance engineering across organizational boundaries. In *Proceedings of the ACM conference on Computer Supported Cooperative Work*, 2002. 2.2.6
- [94] Gunnar Stevens and Volker Wulf. Computer-supported access control. *ACM Trans. Comput.-Hum. Interact.*, 16(3):1–26, 2009. 1, 2.3, 2.3.2, 2.3.3, 2.4
- [95] Oliver Stiemerling and Volker Wulf. Beyond” Yes or No”-Extending Access Control in Groupware with Awareness and Negotiation. *Group Decision and Negotiation*, 9(3):221–235, 2000. 2.2.6, 2.3.3
- [96] Joshua Sunshine, Serge Egelman, Hazim Almuhiemedi, Neha Atri, and Lorrie Faith Cranor. Crying wolf: An empirical study of SSL warning effectiveness. In *Proceedings of the conference on USENIX security symposium*, 2009. 2.1, 3.1.3, 8, 8.1, 9.2.1
- [97] Jennifer Tam, Robert W. Reeder, and Stuart Schechter. I’m allowing what? disclosing the authority applications demand of users as a condition of installation. Technical Report MSR-TR-2010-54, Microsoft, May 2010. 5.3.1, 5.3.2, 9.2.1
- [98] Janice Y. Tsai. *The impact of salient Privacy information on decision-making*. PhD thesis, Carnegie Mellon University, 2009. 12-1-2009. 2.1
- [99] Janice Y. Tsai, Serge Egelman, Lorrie Faith Cranor, and Alessandro Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Journal of Information Systems Research*, 22(2):254–268, 2011. 2.1, 7.2.2
- [100] Maarten W. van Someren, Yvonne F. Barnard, and Jacobijn A.C. Sandberg. *The Think Aloud Method: A practical guide to modelling cognitive processes*. Academic Press, 1994. 7.2.1, 8.6.3
- [101] Kami Vaniea, Clare-Marie Karat, Joshua B. Gross, John Karat, and Carolyn Brodie. Evaluating assistance of natural language policy authoring. In *Proceedings of the Symposium on Usable Privacy and Security*, 2008. 2.3.1, 7.2.2, 9.2.1, 9.2.3
- [102] Kami Vaniea., Lujo Bauer., Lorrie Faith Cranor, M. K. Reiter, and Mike K. Reiter.

- Out of sight, out of mind: Effects of displaying access-control information near the item it controls. In *Proceedings of Privacy Security and Trust*, 2012. 7.2.3, 8.2
- [103] Yang Wang. *A Framework for Privacy-Enhanced Personalization*. Ph.D. dissertation, University of California, Irvine, 2010. 2.1, 8
 - [104] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. I regretted the minute I pressed share: A qualitative study of regrets on Facebook. In *Proceedings of the Symposium on Usable Privacy and Security*, 2011. 1, 1, 7.2.2
 - [105] Ryan West. The psychology of security. *Communications of the ACM*, 51:34–40, 2008. ISSN 0001-0782. 8.1, 8.7
 - [106] Tara Whalen, Diana Smetters, and Elizabeth F. Churchill. User experiences with sharing and access control. In *Proceedings of the extended abstracts on Human Factors in Computing Systems*, 2006. 1, 2.2.1, 2.2.4, 2.3.2
 - [107] Alma Whitten and J. D. Tygar. Why Johnny can’t encrypt: A usability evaluation of PGP 5.0. In *Proceedings of USENIX Security Symposium*, 1999. 8.1, 8.2
 - [108] Michael S. Wogalter. *Communication-Human Information processing (C-HIP) Model*, pages 51–61. Lawrence Erlbaum Associates, 2006. 2.5, 7.2.1, 9.2.1
 - [109] Volker Wulf and Bjorn Golombek. Direct activation: A concept to encourage tailoring activities. *Journal of Behaviour and Information Technology*, 20(4):249–263, 2001. 2.3.2
 - [110] Xia Zhao and M. Erik Johnson. Information governance: Flexibility and control through escalation and incentives. In *Proceedings of the workshop on the Economics of Information Security, Dartmouth College, June*, 2008. 2.3.3

Appendix A

Focus group study

A.1 Focus Group Script

Welcome. We want to thank you for your participation in our focus group on how people share pictures and other documents using the internet. My name is Kami and this is Veda - we are both students at Carnegie Mellon University. I will be moderating today and Veda will be taking notes.

This session will be audio recorded for latter review. Please try and stay on topic during the discussion and try not to say anything that you wouldn't want to be recorded. The topics we will be discussing today should not be of a sensitive nature. However, if at any time you want to say something that you do not want to be recorded please just let me know and I will temporarily turn off the audio recorder.

Your opinions are very important to us, and we want you to feel free to tell us exactly what you think - and we hope, that your ideas will create discussion.

Today we will be talking about sharing documents and pictures online using websites such as Flickr, Facebook, Picasa, YouTube or Myspace. (Icebreaker) To start I want everyone to tell us your first name and a web site you use to share information such as pictures. I'll start, I use a photo sharing software called Gallery to share picutures with friends and co-workers.

I'd like to continue this session with a discussion about your past experiences with sharing electronic files like photographs, music, videos and documents with other people using the computer. I'd like to go around the table again and have everyone tell us about the last time you posted a file to an online sharing site. When I say "file" I mean anything from a Microsoft Word document to a photograph. What were you sharing? Who were you sharing it with and why did you chose that particular way to share it?

If anyone says something interesting ask a question but this section should have limited conversation. Prompts

- *Why did you choose that web site?*
- *I'm less interested in Facebook posts and more interested in Photographs, video or documents such as Word documents.*

Thank you. Now that we have all heard about how the other people at this table share information with people they want to share with. Can anyone tell me about an experience where you discovered that someone you didn't want to see your shared files either could or did see them?

(If no one answers: How likely is it that your shared files can be seen by someone who you don't want to see them?)

Prompts:

- *How did you find out that they saw your files?*
- *Were you able to solve the problem?*
- *Why were they able to see your files?*
- *How did you feel about that person seeing your files?*
- *Do you still feel comfortable sharing files online?*
- *Did you alter how you post files online. For example did you choose to not post some files because of this experience or did you change your privacy settings?*
- *How are you preventing this from happening in the future?*
- *Does anyone else want to share a different experience where sharing files online didn't go as you expected?*

Has anyone had the opposite problem where you tried to share a file with someone and they couldn't see the file?

Prompts:

- *I'm less interested in email and technical issues and more interested in situations where your settings prevented them from seeing the file. For example if I shared pictures on Facebook and I only wanted my friends to see it not my Mom so I only shared with friends and latter realized that my sister, who I wanted to see the pictures couldn't see them.*
- *How did you find out that they couldn't see your files?*
- *Were you able to solve the problem?*
- *Why couldn't they see your files?*
- *How are you preventing this from happening in the future?*

Hand out comics.

Now I would like to move on. You talked about sharing information using *[insert example from prior conversation]*. Now imagine a photo sharing web site had a feature where you could see who has been looking at your shared photos and who could look at your photos. I've handed you comics about two people named Alice and Joe who use a web site like this. Please read their stories.

Can you imagine an instance where you or a friend might experience a situation like those Alice and Joe encountered?

Prompts:

- *Can you see yourself or a friend using information about who has seen your pictures to reconnect with a friend?*
- *Can you see yourself or a friend using information about who could see your pictures to identify people who can see your pictures but shouldn't?*

I'm now going to give each of you a packet with some example photo sharing websites.

We are going to go through each page of the packet together so please do not look ahead in the packet.

Hand out packets

Please open the packet like this (*demonstrate opening so both the websites are visible.*) so you can see two pages at once.

The first two pages of the packet are screen shots of a potential photo sharing web site that lets you organize and share your photos. I'd like you to imagine that this is your favorite photo sharing website and that it already has all the features you are used to seeing. If I click on one of the photo albums it will open and show the pictures inside.

It has a feature where the owner of a photo album can see information about who has and and who could see their photographs. We are going to use the projector to show you how this website might work. Your comments and opinions are extremely important to us so feel free to write on any of the pages in the packet including the pictures. I'm going to collect the packets at the end so if there is anything you thought was important but didn't get to say please write it down.

In this first example (*describe the interface*)

Allow participants to make comments at this point. If they ask questions about things covered in the written description answer them, if not ask the participant what they think it would look like or what they think it should do.

On the next page there are several questions about this website. Its important to remember we are testing our sample website designs, our vocabulary and layout choices not you.

The questions on this page are designed to represent several different questions people might try to answer if they had information about who could see their pictures and who has seen their pictures. They are supposed to assist you in understanding how you might use this webpage so you can give more informed opinions about it as well as compare it to the other websites I will be showing you. Not all the questions can be answered and some have ambiguous answers. If you feel that it is impossible to answer a question just write down that it can't be answered. We are testing the webpage layout not you. There are no wrong answers to these questions. Also, if anything seems particularly confusing about the website design I would like you to circle it so we can discuss it latter.

Do you have any questions?

Please try and answer the questions on your own right now.

Wait for the majority to answer the questions

I'd like to move on to a discussion of this website design now. Its all right if you haven't finished answering the questions. Feel free to write any additional comments you have during the discussion. After interacting with this interface do you think it is something you would like to use as part of your favorite photot sharing website?

[use prompts below]

For each pair of information display pages in the packet repeat the following script.

Please turn to the next page. (*describe the interface*)

Allow participants to make comments at this point. If they ask questions about things covered in the written description answer them, if not ask the participant what they think it would look like or what they think it should do.

Please look at the website and answer the questions in the provided space on the second page. Feel free to draw on the webpage screenshot and point out anything you think is confusing. I'll give you a few minutes to do so.

Wait for the majority to answer the questions

Now that everyone has looked at the website can someone start us out by saying what they think the best and worst thing about this website is?

Prompts:

- *Would this website be useful for Alice?*
- *Would this website be useful for Joe?*
- *What do you think the feature shown in this webpage would be useful for?*
- *Was any of the language on this page confusing?*
- *If you saw this information display next week how confident are you that you could use it?*
- *What did you like or find confusing?*
- *Which was more useful in this interface: who could see the pictures or who did see the pictures?*
- *If you could change the way the website looks, what would you change?*

After going through the whole packet

Now that you have seen several different ways of showing information about who has and who could see pictures in an online photo album, I'd like to go around the table and have each person say what their favorite and least favorite website was and why.

Prompts:

- *Of the different types of information you have seen presented today which do you find to be the most useful?*
- *Are you more interested in who looked, what was looked at or how often it happened?*

I would like to thank everyone for coming. Please leave your packets on the table.

A.1.1 Information visualization explanations

Website A Information about who has and who could see each of these albums is listed below the album name. For example Alex, Jane and four other people, who's names are not listed, have seen Halloween 2009 photos. A total of six people have the ability to view the album.

- Would you prefer to see who the "potential viewers are"?
- What else might you want to find out about your photo use that this application isn't showing you?

Website B This webpage shows information about who has and who could view the albums shown on this page as well as anything inside those albums.

On the left side of this webpage there is a grid of people across the top and albums down the left side. The colors in the grid indicate if that person can see that album, green means they can see anything in the album, yellow means they can see some pictures in the album but not all and red means they can't see anything.

The numbers indicate how often they have looked at the album. At the bottom left there is a small bar graph showing how often people have looked at any of the albums

over a long time. The small window is showing the time period where the numbers are coming from.

For example if I were to look at Nicole I can see that she can see the Niagara Falls pictures and that she has looked at one picture in the album in the last three months. If I select Nicole the albums she could see are all highlighted. The highlight color indicates how often Nichole looked at the pictures in that album. Dark blue means the most and light blue means the least.

- Do the colored frames around the pictures make sense?
- Can anyone tell me what the bar graph in the bottom left means?
- Why is “Around Pittsburgh” colored yellow? What does that mean?

Website C On the top of this webpage there is a list of people and a graphic showing information about who could view these albums. The list of people on the left shows who has been looking at pictures. People above the dotted line have looked at some of Alice’s pictures in the last month.

Albums at this website can contain other albums inside of them. For example “Around Pittsburgh” may contain another album called “The Strip.” The graphic shows all the albums including some of the albums inside of other albums.

The graphic also shows what albums the highlighted person has or could see. If an album is green than the highlighted person can see anything in that album. If the color is yellow then that person can see some of the pictures in that album and red indicates that the person can see nothing in that album.

The bigger the rectangle that represents the album the more times that person has looked at that album. If I were to click on one of the names the graphic would change to show what albums that user has and could see.

- This website shows you the policies of every album and subalbum that instead of just the albums in this folder. Is this useful to you?
- The list on the left shows at a glance who has been recently looking at photographs. Is a name with no context sufficient to understand what is going on.

Website D In this website information about who can and has seen pictures in any of the albums is shown on the left. There is a list of people in this box. On the left of each person’s name is a colored box, if it is green they can see any of the albums, yellow means they can see some of the albums and red means they can’t see any albums.

The small graph to the right of the person’s name shows when they saw pictures. The start and end dates for this graph are indicated by the labels on the top. In this case they go from January to April.

- Would you think to click on the names on the left to determine what they looked at?
- Is it clear how long the graph next to the names is for?
- Is it easy to understand the re-sizing of the images?

- Is it easy to understand why some of the album pictures are greyed out?

Website E Information about who has and who could see each of these albums is listed below the album name. For example the Phipps photos were seen by 6 people and could be seen by Alice, Kate and 6 other people. The small graph indicates when the album was viewed over the last month.

- Can anyone tell me one person who recently viewed the Phipps photos? Is it clear that the names are people who could view not people who have viewed?
- Do the small graphs make sense?
- Do you think you would casually look at this information when viewing your online photo albums?

Website F On the left of this application are several sections each labeled with a person's name. Below each label are several albums that person has seen over the last month. The bigger the name the more often they saw the album. Black albums have been seen recently and they fade to grey as time passes. After a month they completely disappear.

- Is a month long enough?
- Is it clear what the names of the albums are in the information display?
- Would you expect to see names of subalbums here?

Website G On the left of this application is a list of the people who have and can see any of the albums. "Who has seen my pictures" is ordered starting with the person who most recently viewed an album. Next to each name is the last time they saw a picture and how long they looked at the pictures on that occasion. Below is a list of all the people who can see at least one picture in these albums.

- Is the length of time they looked at your albums interesting?
- Is the list of who could see pictures interesting even though you don't know what they can see?

A.2 Focus group 1 packet

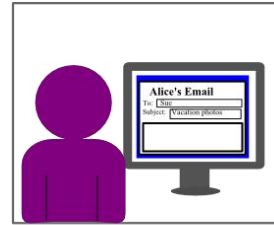


Alice is looking through her online photographs for ones she can use for her screensaver.



Sue recently looked at this picture.

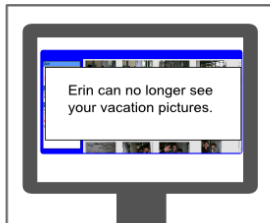
Alice is delighted to discover that her good friend Sue found time to look at the pictures of their trip to Chicago together.



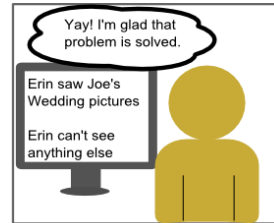
Alice hasn't had a chance to talk to Sue much since they got back so she decides this is the perfect time to reconnect. Alice writes Sue an email saying how much fun she had and she can't wait to see Sue again at Christmas.



Joe is posting pictures from his vacation on his favorite photo sharing website. He is looking through the pictures to make sure they are all ok when he notices that his ex-girlfriend, Erin, can still see his vacation pictures.

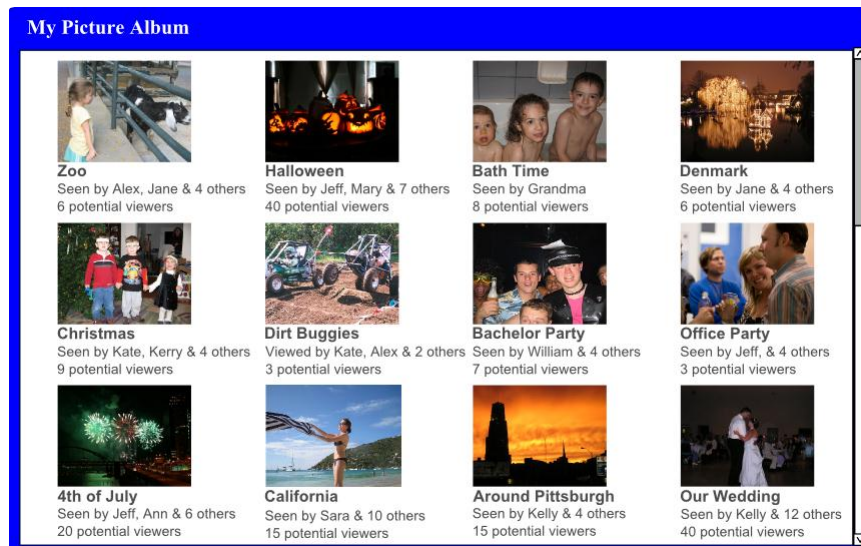


Joe is very upset and quickly removes Erin from the list of people who can see his vacation pictures.



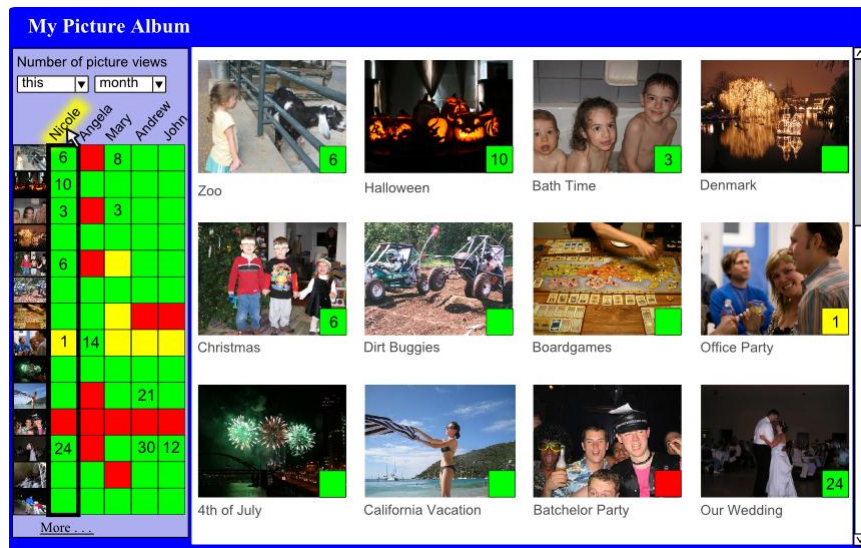
Joe is very concerned that Erin may have seen some pictures he doesn't want her to see. So he looks through the remainder of his photo album to make sure she can't see anything else and to check that she hasn't seen anything she shouldn't.

Website A



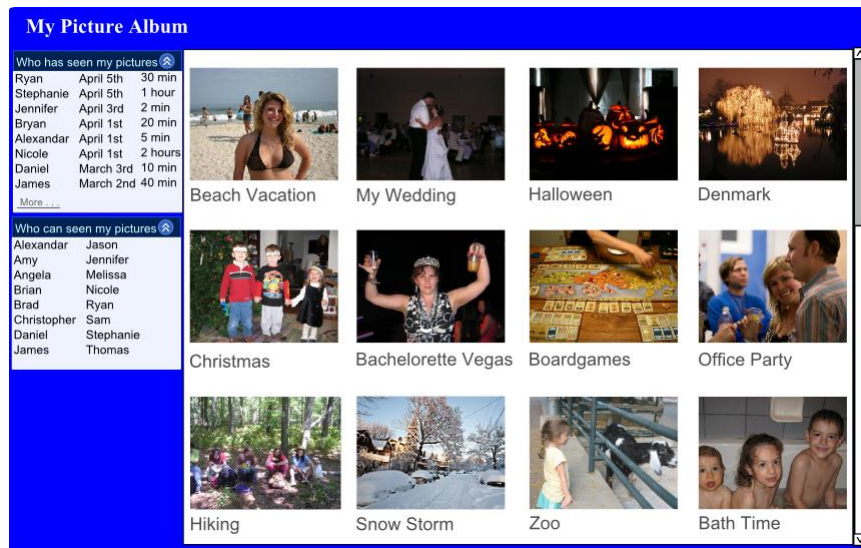
1. Name a person who can see the "Around Pittsburgh" pictures.
2. Name one album Jeff can see.
3. Name one album that Sara has viewed.
4. Name one person who has viewed the Denmark photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album that was viewed today?
8. Name a person who viewed an album today?
9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B



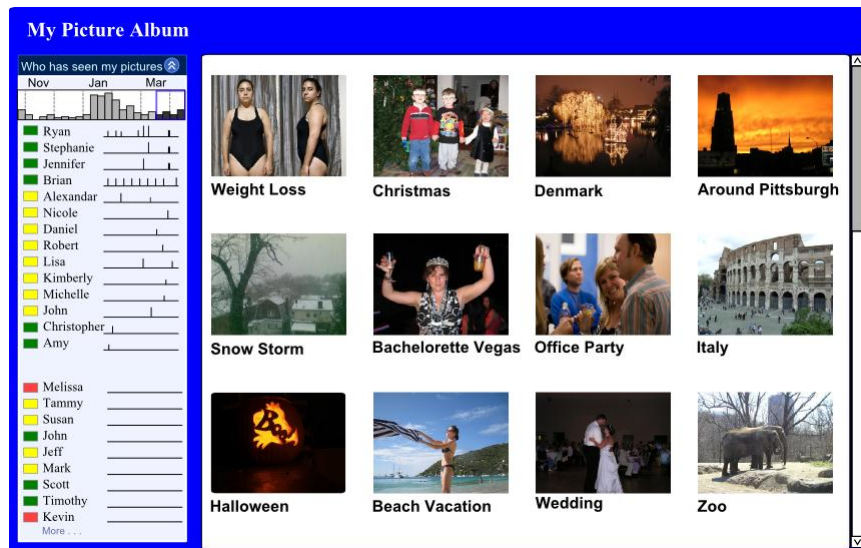
1. Name a person who can see the "4th of July" pictures.
2. Name one album Mary can see.
3. Name one album that Andrew has viewed.
4. Name one person who has viewed the Boardgames photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed this month.
8. Name a person who viewed an album this week.
9. Nicole claimed that she looked at "Our Wedding" pictures over Christmas vacation. Is this true?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C



1. Name a person who can see the "Snow Storm" pictures.
2. Name one album James can see.
3. Name one album that Alexandar has viewed.
4. Name one person who has viewed the Denmark photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed in March?
8. Name a person who viewed an album in the first week of April?
9. Of the people who recently looked at pictures who spent the most time?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D



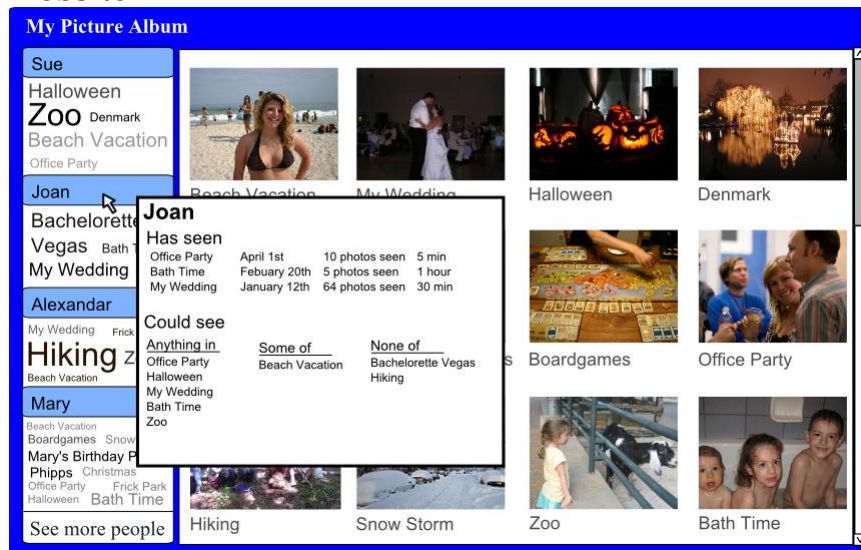
1. Name a person who can see the "Weight Loss" pictures.
2. Name one album that Nicole can see.
3. Name one album that Stephanie has viewed.
4. Name one person who has viewed the "Beach Vacation" photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. What is the most recently viewed album?
8. Who was the last person to view a picture?
9. Which friend views your pictures regularly?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E



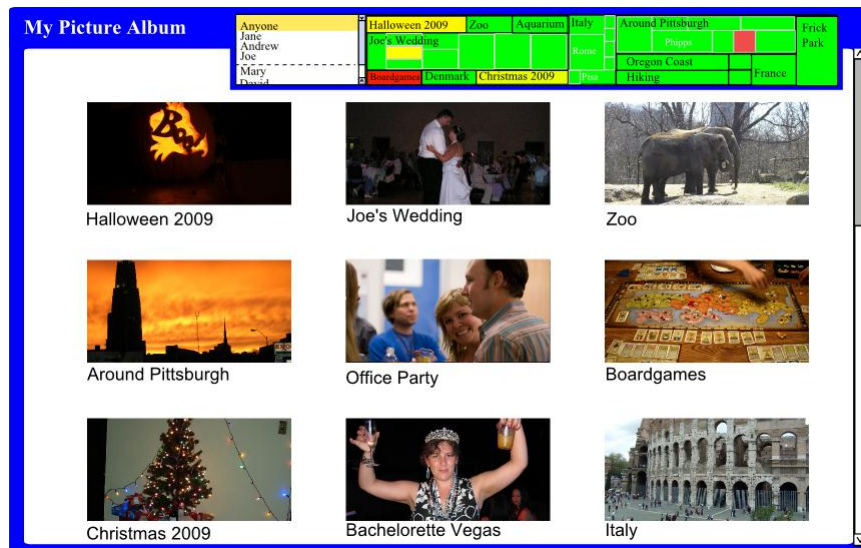
1. Name a person who can see the "Frick" pictures.
2. Name one album Brian can see.
3. Name one album that Jennifer has viewed.
4. Name one person who has viewed the Christmas 2009 photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed this month.
8. Name a person who has looked at an album this month.
9. You recently emailed out a link to one of her albums to a large number of her friends. Which album was it?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F



1. Name a person who can see the "Hiking" pictures.
2. Name one album Sue can see.
3. Name one album that Mary has viewed.
4. Name one person who has viewed the Zoo photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album that was recently viewed.
8. Name a person who recently viewed an album.
9. Which of your friends likes to glance at lots of your pictures?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

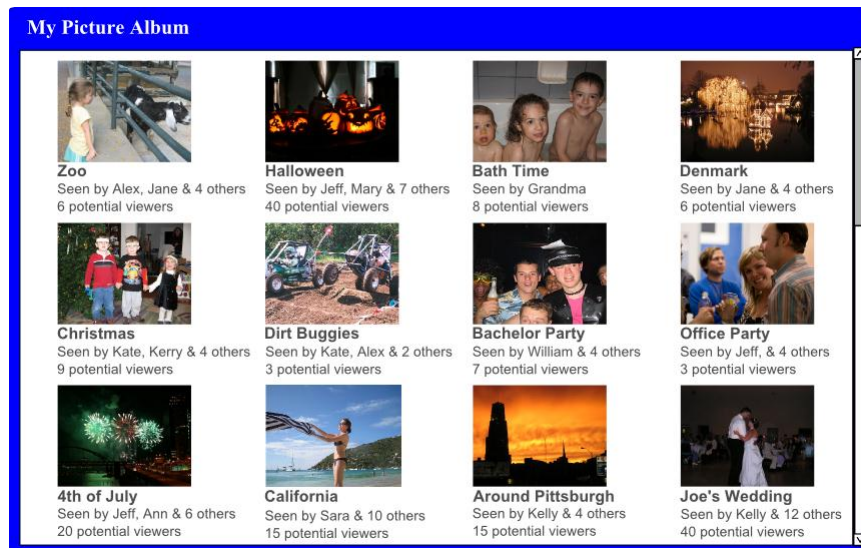
Website G



1. Name a person who can see the "Joe's Wedding" pictures.
2. Name one album Joe can see.
3. Name one album that Andrew has viewed.
4. Name one person who has viewed the Italy photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. What is the most recently viewed album?
8. Who was the last person to view a picture?
9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

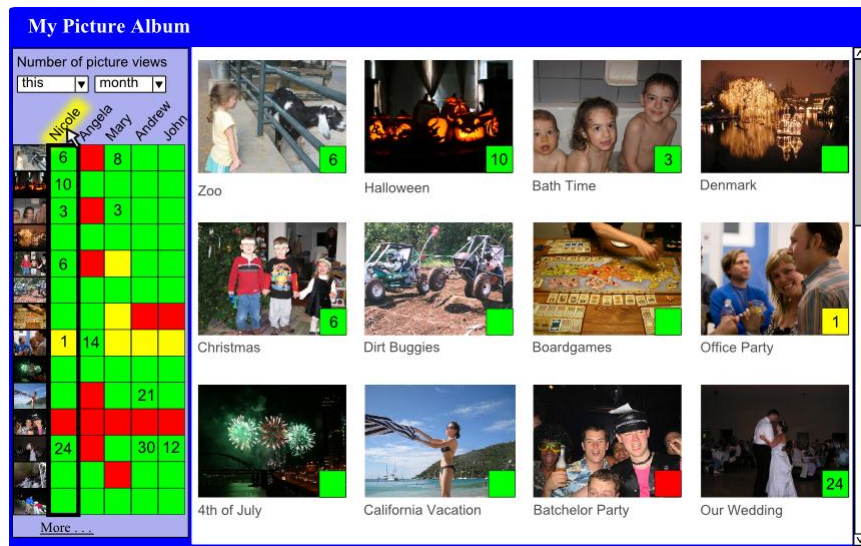
A.3 Focus group 2 packet

Website A



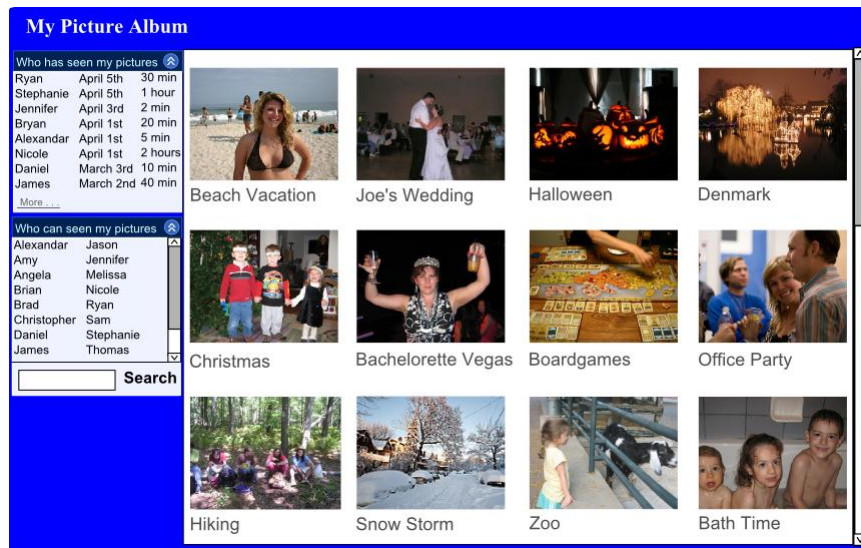
1. Name a person who can see the "Around Pittsburgh" pictures.
2. Name one album Jeff can see.
3. Name one album that Sara has viewed.
4. Name one person who has viewed the Denmark photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album that was viewed today?
8. Name a person who viewed an album today?
9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B



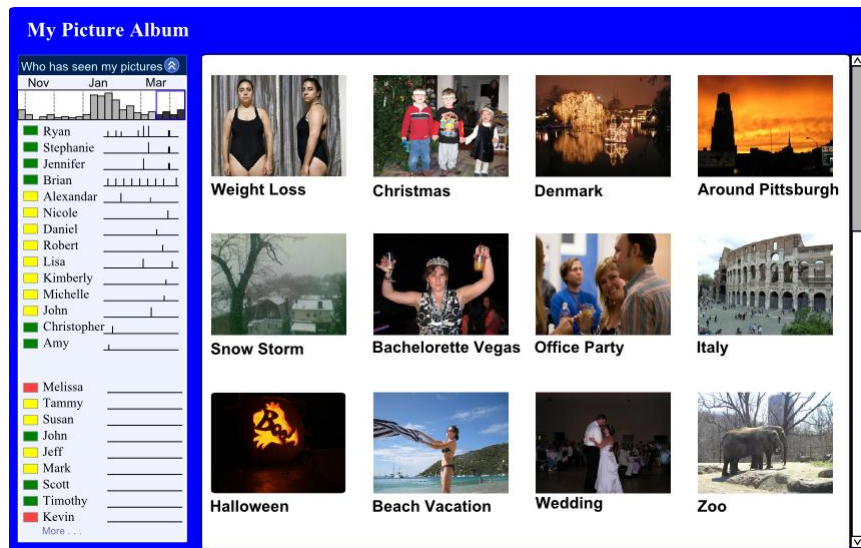
1. Name a person who can see the "4th of July" pictures.
2. Name one album Mary can see.
3. Name one album that Andrew has viewed.
4. Name one person who has viewed the Boardgames photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed this month.
8. Name a person who viewed an album this week.
9. Nicole claimed that she looked at "Our Wedding" pictures over Christmas vacation. Is this true?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C



1. Name a person who can see the "Snow Storm" pictures.
2. Name one album James might be able to see.
3. Name one album that Alexandar has viewed.
4. Name one person who has viewed the Denmark photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed in March?
8. Name a person who viewed an album in the first week of April?
9. Of the people who recently looked at pictures who spent the most time?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D



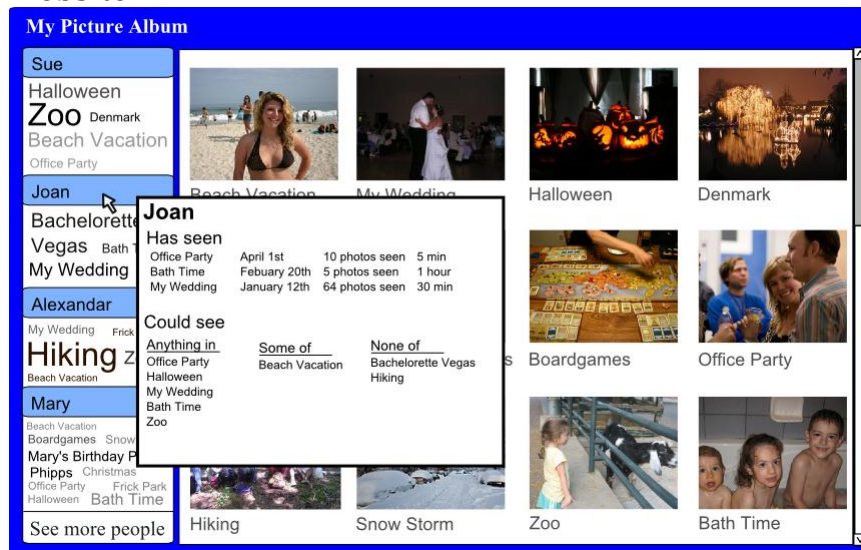
1. Name a person who can see the "Weight Loss" pictures.
2. Name one album that Nicole can see.
3. Name one album that Stephanie has viewed.
4. Name one person who has viewed the "Beach Vacation" photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. What is the most recently viewed album?
8. Who was the last person to view a picture?
9. Which friend views your pictures regularly?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E



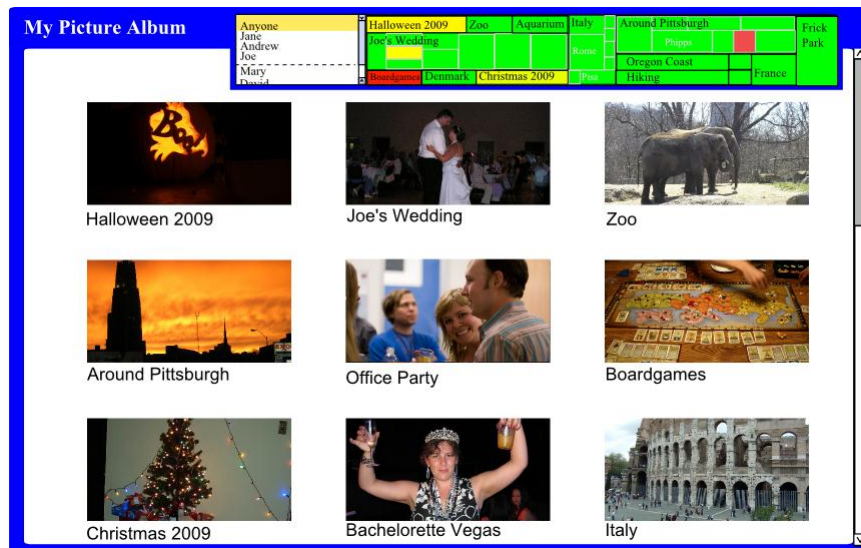
1. Name a person who can see the "Frick" pictures.
2. Name one album Brian can see.
3. Name one album that Jennifer has viewed.
4. Name one person who has viewed the Christmas 2009 photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed this month.
8. Name a person who has looked at an album this month.
9. You recently emailed out a link to one of her albums to a large number of your friends. Which album was it?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F



1. Name a person who can see the "Hiking" pictures.
2. Name one album Sue can see.
3. Name one album that Mary has viewed.
4. Name one person who has viewed the Zoo photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album that was recently viewed.
8. Name a person who recently viewed an album.
9. Which of your friends likes to glance at lots of your pictures?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website G



1. Name a person who can see the "Joe's Wedding" pictures.
2. Name one album Jason can see.
3. Name one album that Alexandar has viewed.
4. Name one person who has viewed the Christmas photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. What is the most recently viewed album?
8. Who was the last person to view a picture?
9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

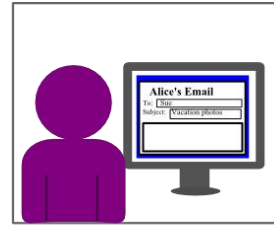
A.4 Focus group 3 packet



Alice is looking through her online photographs for ones she can use for her screensaver.



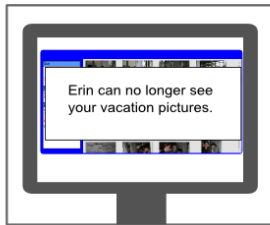
Alice is delighted to discover that her good friend Sue found time to look at the pictures of their trip to Chicago together.



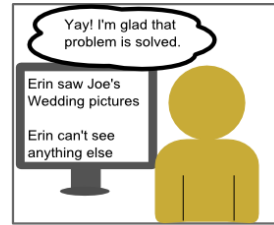
Alice hasn't had a chance to talk to Sue much since they got back so she decides this is the perfect time to reconnect. Alice writes Sue an email saying how much fun she had and she can't wait to see Sue again at Christmas.



Joe is posting pictures from his vacation on his favorite photo sharing website. He is looking through the pictures to make sure they are all ok when he notices that his ex-girlfriend, Erin, can still see his vacation pictures.

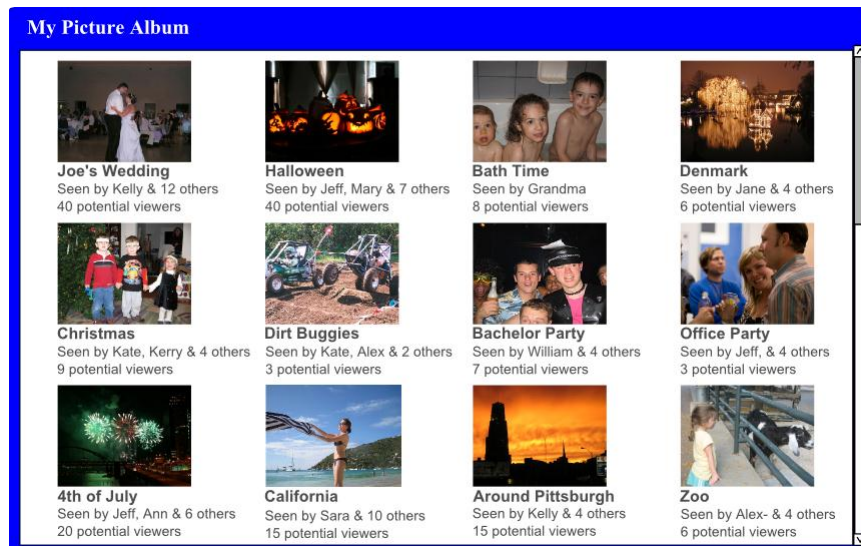


Joe is very upset and quickly removes Erin from the list of people who can see his vacation pictures.



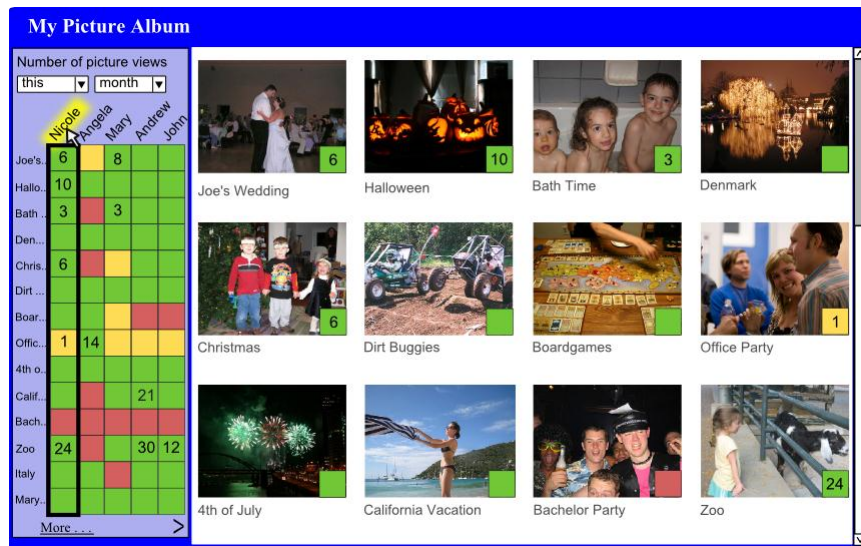
Joe is very concerned that Erin may have seen some pictures he doesn't want her to see. So he looks through the remainder of his photo album to make sure she can't see anything else and to check that she hasn't seen anything she shouldn't.

Website A



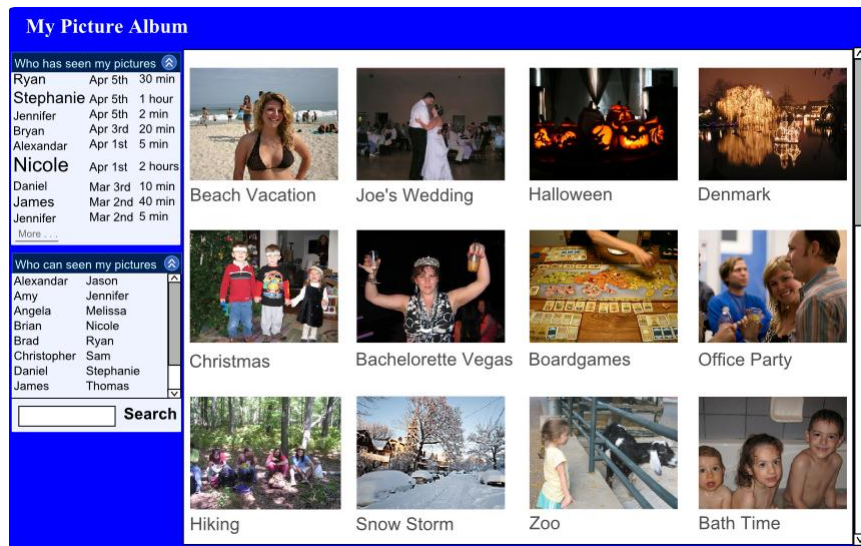
1. Name a person who can see the "Around Pittsburgh" pictures.
2. Name one album Jeff can see.
3. Name one album that Sara has viewed.
4. Name one person who has viewed the Denmark photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album that was viewed today?
8. Name a person who viewed an album today?
9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B



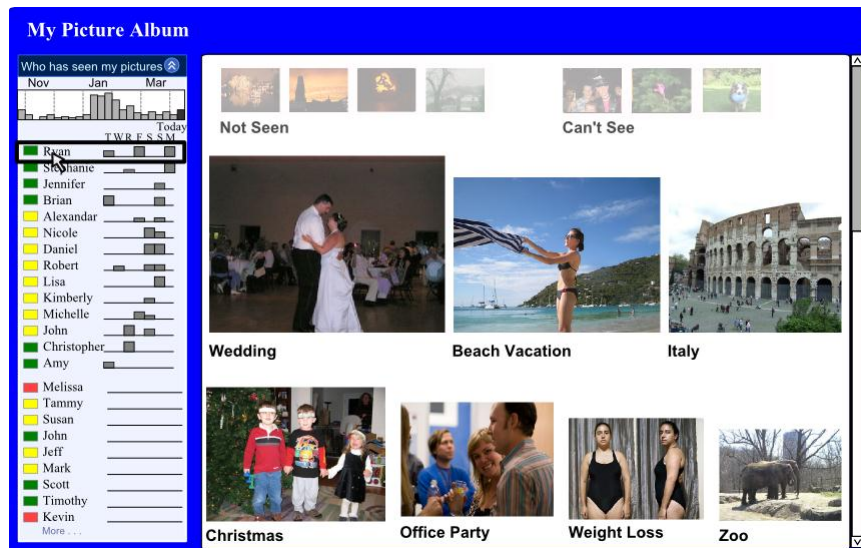
1. Name a person who can see the "4th of July" pictures.
2. Name one album Mary can see.
3. Name one album that Andrew has viewed.
4. Name one person who has viewed the Boardgames photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed this month.
8. Name a person who viewed an album this week.
9. Nicole claimed that she looked at "Our Wedding" pictures over Christmas vacation. Is this true?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C



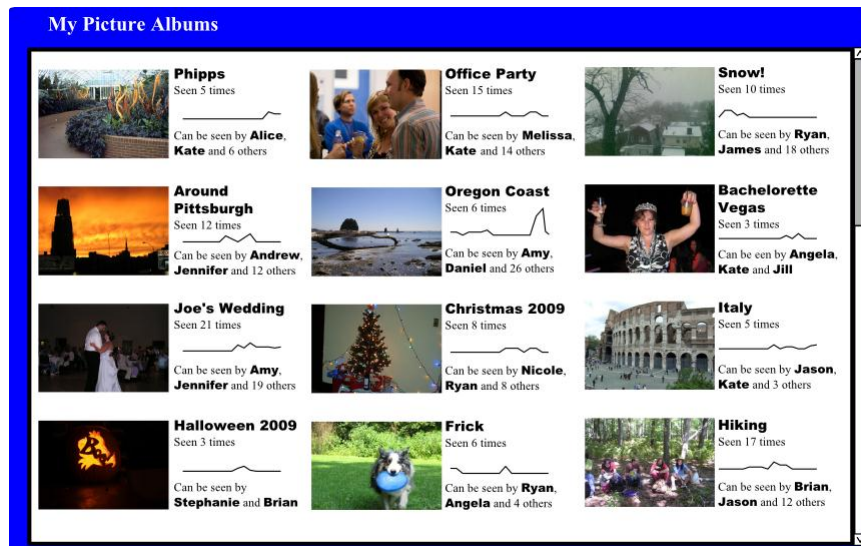
1. Name a person who can see the "Snow Storm" pictures.
2. Name one album James might be able to see.
3. Name one album that Alexandar has viewed.
4. Name one person who has viewed the Denmark photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed in March?
8. Name a person who viewed an album in the first week of April?
9. Of the people who recently looked at pictures who spent the most time?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D



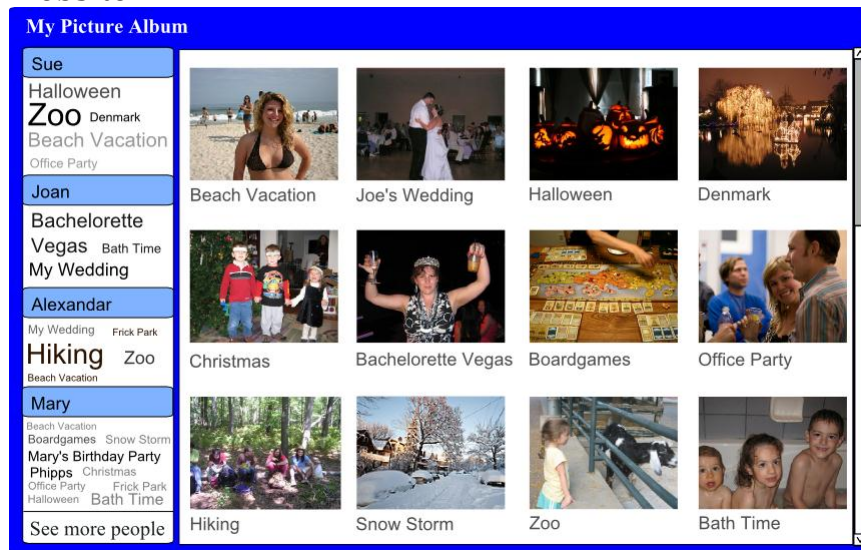
1. Name a person who can see the "Weight Loss" pictures.
2. Name one album that Nicole can see.
3. Name one album that Stephanie has viewed.
4. Name one person who has viewed the "Beach Vacation" photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. What is the most recently viewed album?
8. Who was the last person to view a picture?
9. Which friend views your pictures regularly?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E



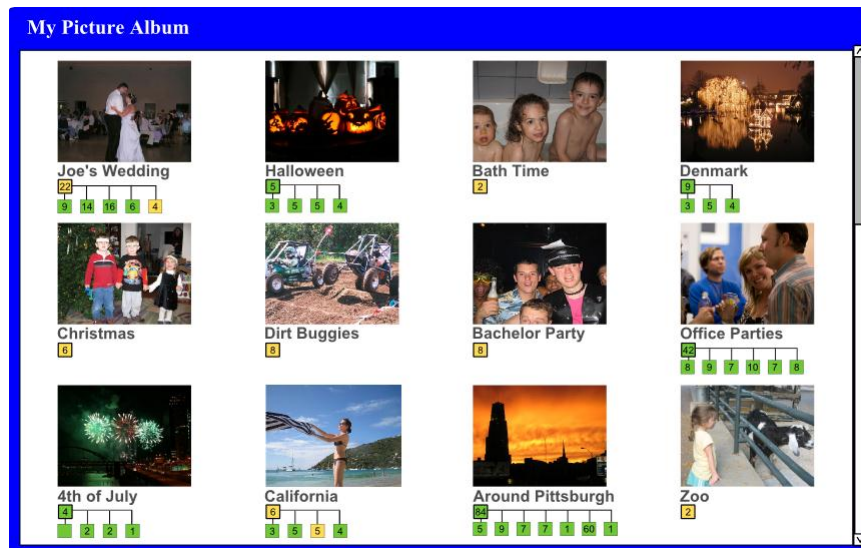
1. Name a person who can see the "Frick" pictures.
2. Name one album Brian can see.
3. Name one album that Jennifer has viewed.
4. Name one person who has viewed the Christmas 2009 photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album viewed this month.
8. Name a person who has looked at an album this month.
9. You recently emailed out a link to one of her albums to a large number of your friends. Which album was it?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F



1. Name a person who can see the "Hiking" pictures.
2. Name one album Sue can see.
3. Name one album that Mary has viewed.
4. Name one person who has viewed the Zoo photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. Name an album that was recently viewed.
8. Name a person who recently viewed an album.
9. Which of your friends likes to glance at lots of your pictures?
10. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website G



1. Name a person who can see the "Joe's Wedding" pictures.
2. Name one album Jason can see.
3. Name one album that Alexandar has viewed.
4. Name one person who has viewed the Christmas photos.
5. What is the most frequently viewed album?
6. Who viewed the most photos?
7. What is the most recently viewed album?
8. Who was the last person to view a picture?
9. What do you like or find confusing about this interface? Feel free to draw on the website picture.

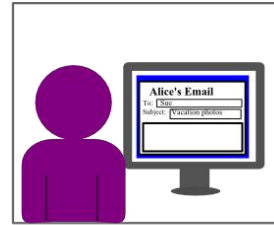
A.5 Focus group 4 and 5 packet



Alice is looking through her online photographs for ones she can use for her screensaver.



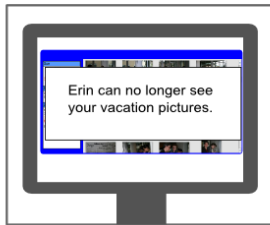
Alice is delighted to discover that her good friend Sue found time to look at the pictures of their trip to Chicago together.



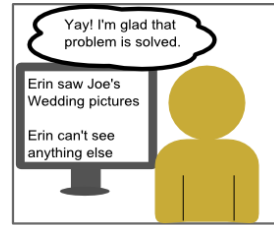
Alice hasn't had a chance to talk to Sue much since they got back so she decides this is the perfect time to reconnect. Alice writes Sue an email saying how much fun she had and she can't wait to see Sue again at Christmas.



Joe is posting pictures from his vacation on his favorite photo sharing website. He is looking through the pictures to make sure they are all ok when he notices that his ex-girlfriend, Erin, can still see his vacation pictures.

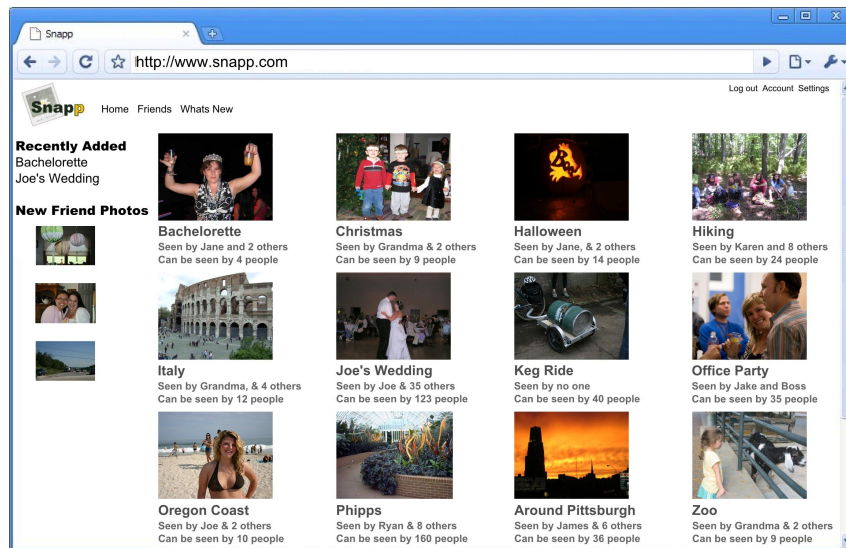


Joe is very upset and quickly removes Erin from the list of people who can see his vacation pictures.



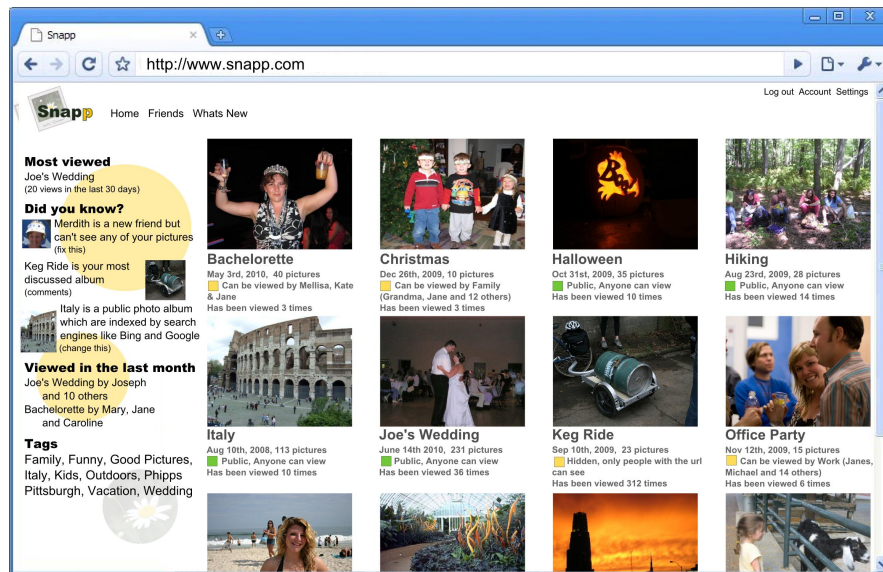
Joe is very concerned that Erin may have seen some pictures he doesn't want her to see. So he looks through the remainder of his photo album to make sure she can't see anything else and to check that she hasn't seen anything she shouldn't.

Website A



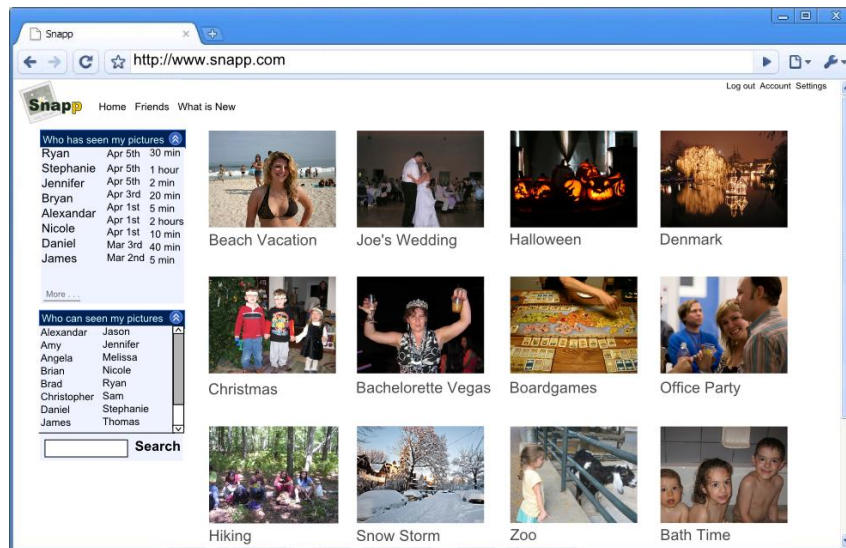
1. Name a person who can see the "Around Pittsburgh" pictures.
2. Name one album Grandma has seen.
3. What is the most frequently viewed album?
4. Name an album that was viewed today?
5. Would this website have helped Alice? Why?
6. Would this website have helped Joe? Why?
7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website B



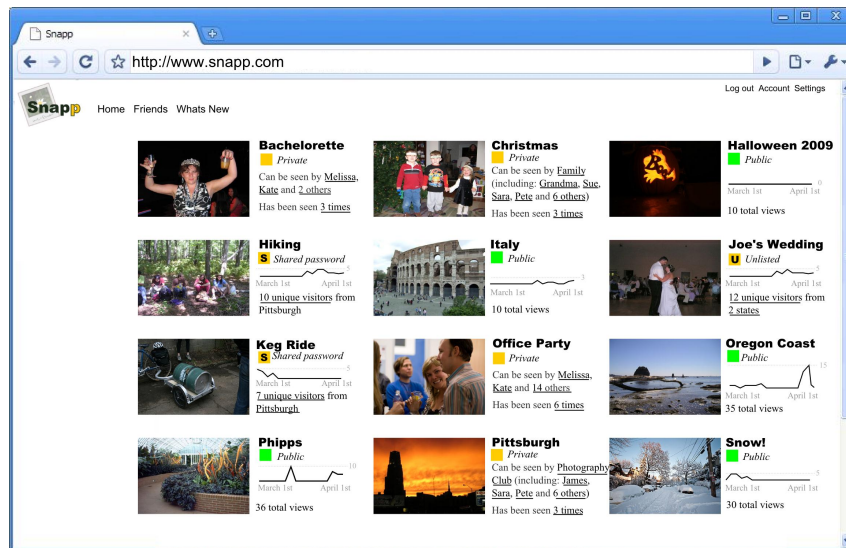
1. Who can see Joe's Wedding pictures.
2. Name one album Merdith **cannot** see.
3. What is the most frequently viewed album?
4. Name an album that was viewed today?
5. Would this website have helped Alice? Why?
6. Would this website have helped Joe? Why?
7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website C



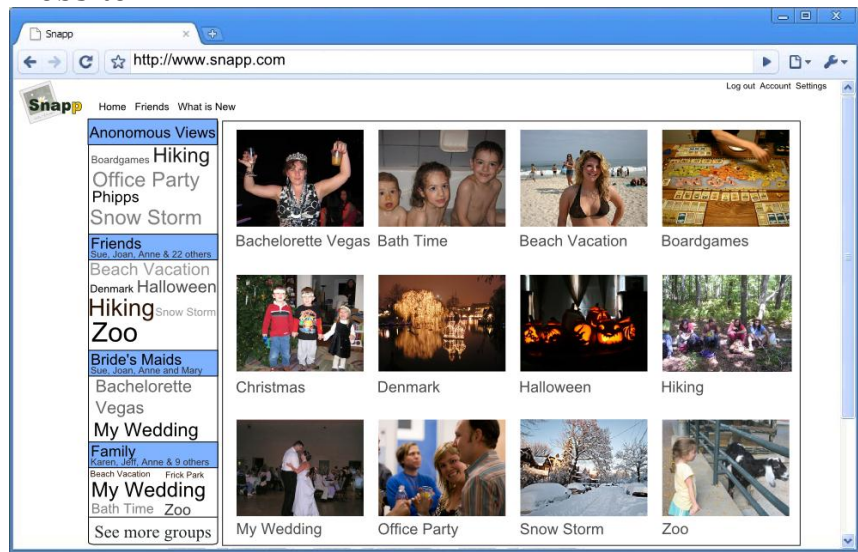
1. Name a person who can see the "Snow Storm" pictures.
2. Name one album Alexandar can view.
3. What is the most frequently viewed album?
4. Name an album that was viewed today?
5. Would this website have helped Alice? Why?
6. Would this website have helped Joe? Why?
7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website D



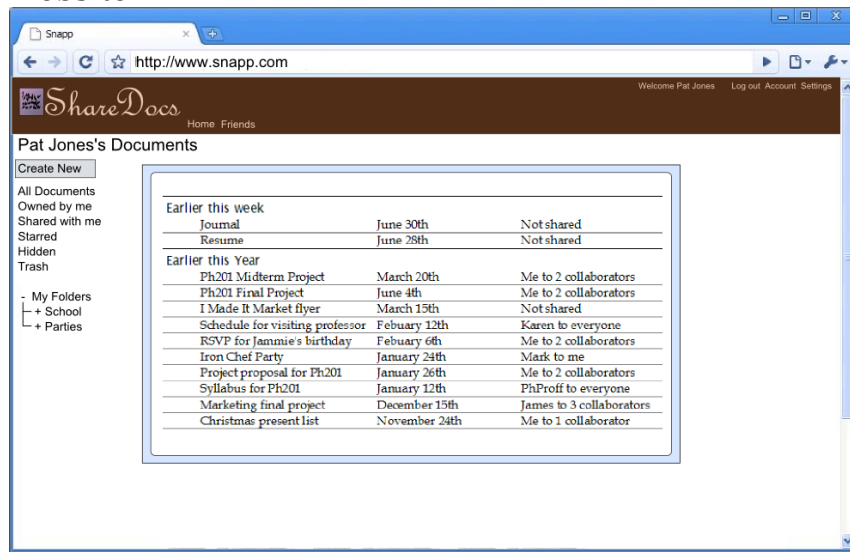
1. Name a person who can see the Office Party pictures.
2. Name one album that Nicole can see.
3. What is the most frequently viewed album?
4. Name an album that was viewed today?
5. Would this website have helped Alice? Why?
6. Would this website have helped Joe? Why?
7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website E



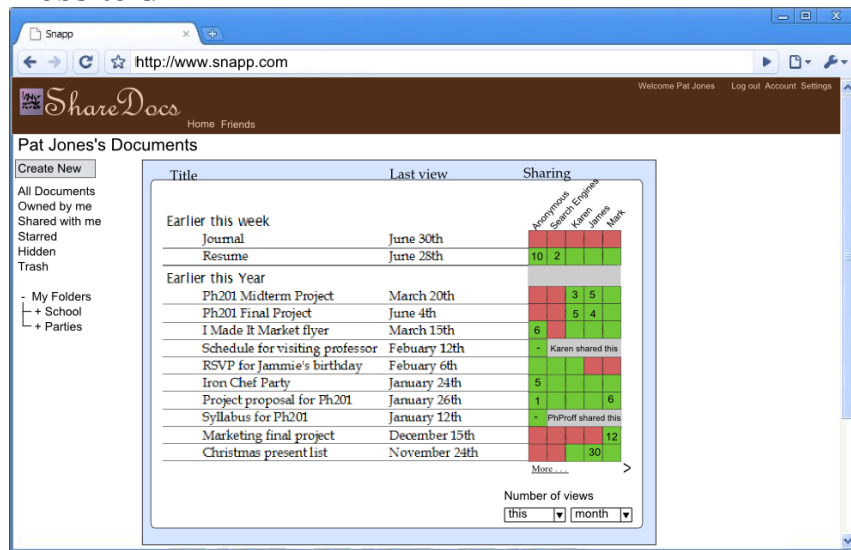
1. Name a person who can see the "Hiking" pictures.
2. Name one album Sue can see.
3. What is the most frequently viewed album?
4. Name an album that was viewed today?
5. Would this website have helped Alice? Why?
6. Would this website have helped Joe? Why?
7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website F



1. Name a person who can see the "Syllabus for Ph201".
2. Name one document Karen can see
3. Name one document created by a friend.
4. Name a document last seen this week.
5. Would this website have helped Alice?
6. Would this website have helped Joe?
7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Website G



1. Name a person who can see the Ph201 Midterm Project.
2. Name one document James can see.
3. What is the most frequently viewed document?
4. Name a document that was viewed this week?
5. Would this website have helped Alice? Why?
6. Would this website have helped Joe? Why?
7. What do you like or find confusing about this interface? Feel free to draw on the website picture.

Appendix B

Eye-tracker study (study 2)

B.1 Printed instructions and emails

In the eye tracker study (study 2) participants were given instructions and emails by the researcher on printed sheets of paper. The remainder of this appendix section is all the instructions and emails used in the study. Each instruction or email was printed on its own sheet of paper, but in the interests of saving space, we show only the textual content of the pages. Each box of text was printed on a single page, without the black border.

The pages which give instructions and the pages with emails that initiate tasks were given to all participants. The pages with emails used to prompt the participant, were given to the participant only if the participant did not complete all parts of the task.

Instructions

In this study you will be asked to role play a person called Pat Jones. Every time you have to make a decision or judgment call, I want you to think about how Pat Jones would handle the situation and handle it that way.

During this study you should think about the photo albums you are working with as your own (well, Pat Jones's). If you see something that you would change in your own album then go ahead and change it or just say it out loud so I know what you would have changed if you had time.

Today I will give you several information pages and emails written on pieces of paper. Some of the emails will contain simple and straightforward tasks and some will be less directed to get a better sense of how you approach and complete photograph management tasks in general. When you are ready to respond to an email just say out loud what you would email back. Once you have responded I will hand you another piece of paper with the next email.

We are interested in how you approach and solve the issues presented to you. Remember, we are testing the software and how it supports how you work with photographs. We are not testing you.

Say "Done" when you are finished reading this page.

Instructions: Pat Jones

Your name is Pat Jones. You are an administrative assistant at a large web hosting and data storage company called Global Storage. Your company uses a popular online photo sharing site called Gallery to store and share their photographs. You, your family and most of your friends also use Gallery to store and share photographs.

You use Gallery because it gives each person lots of space, it makes it easy to share with only certain groups of people and it lets people, like your Mom, give others the ability to administer their albums for them without having to give out the password. This makes it easier to help your friends and family when they have problems.

Global Storage has a company wide album on Gallery where company related photographs are posted. As an administrative assistant at Global Storage, one of your jobs is to take photographs of events and post them in the company album. The last administrative assistant wasn't very good at this and left errors all through the albums which you clean up as you find them. Your boss and coworkers often ask you to do photo management tasks to keep the company photo album in order and looking good.

All the Global Storage photographs are in the album called "Global Storage" though some employees keep photographs in their personal albums.

Say "Done" when you are finished reading this page.

To: Pat Jones <pat@globalstorage.com>
From: Angela Wilson <angela@globalstorage.com>
Subject: Sideways photograph

Hello Pat,

I was looking through the *Around the office* album in the "Global Storage Shared Albums" and I noticed that Gerald's photograph is sideways.

Could you please fix that.

Thanks,
Angela

To: Pat Jones <pat@globalstorage.com>
From: Angela Wilson <angela@globalstorage.com>
Subject: Sideways photograph

Hello Pat,

Gerald's photograph still appears to be sideways. You can find it if you go into the "Global Storage Album" and then go to "Around the Office". Gerald's photograph is in the upper right hand corner.

Thanks,
Angela

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I noticed that there is an album entitled *Around the office* inside of the *Global Storage Shared Album* album. The photographs you have there are really great! If I email a link to someone at another company, will they be able to see the photos in that album? Its ok if they can't I just want to know before I send an email.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Emailing photographs

Hi Pat,

Are you sure my friends who aren't in the company will be able to see the photos? I remember doing this before and it didn't work . . .

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Emailing photographs

Hi Pat,

Are you sure? I was about to send off the email when Angela dropped by and she swears she saw you looking at the wrong album when you emailed me.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Kevin Brown <kevin@globalstorage.com>
Subject: Remove photos of me

Hi Pat,

I heard this horrible rumor that you put all our photographs on the Internet where anybody could see them and now the boss is emailing the photos to his friends? I know you take great photographs but I look horrible in photos and I really don't want that on the Internet. Could you please delete the photo of me in the People album?

Thanks,
Kevin

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Remove photos of me

Hi Pat,

Kevin really wants his photo taken down. I had a bit of a talk with him about it because I think it is important to have these photos up. The compromise was that you would take the photograph down and I would have our professional photographer take a photograph of Kevin and put it up later.

So please remove Kevin's photograph from the People album.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I checked with human resources and our lawyer, it is fine to allow employees to add photos to an online album. So go ahead and give Global Storage employees (coworkers) the ability to add photos to the “Around the office” album.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I’ve been going through the company photo albums all afternoon. Great new photos by the way. I noticed that there is an album called *Around the office* which seems to be great set of photos of day-to-day events in the office. I’ve noticed that other people sometimes take photographs around the office but they don’t seem to be able to add them to this album.

I’d love it if you made it so other people in the office could add to the *Around the office* album. That way we can have all these great pictures in one place.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Emailing photographs

Hi Pat,

I know you said you fixed it so other Global Storage employees could add to the “Around the office” album but I just tried and it didn’t work. Could you fix it?

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Steve Johnson <steve@globalstorage.com>
Subject: Product fair titling issues

Hello Pat,

The boss has finally decided to pay attention to your photographs; he has been sending me pictures of myself on and off all day. The last one he sent was from the Project Fair and I noticed that you had mistitled my poster. Actually, it looks like you may have switched my title with someone else's, so theirs is wrong too. You should be able to get the correct titles by reading the posters behind each person.

Thanks,
Steve

To: Pat Jones <pat@globalstorage.com>
From: Steve Johnson <steve@globalstorage.com>
Subject: RE: Product fair titling issues

You can find the titles of the posters by looking at the photographs. You can easily read each title behind the person if you just open the photograph instead of looking at the thumbnail.

Sorry, I can't remember my exact title right now.

Thanks,
Steve

To: Pat Jones <pat@globalstorage.com>
From: Steve Johnson <steve@globalstorage.com>
Subject: Product fair titling issues

Hello Pat,

Um, I noticed that you fixed one of the poster titles but not the other one. Could you go fix the other title please?

Thanks,
Steve

To: Pat Jones <pat@globalstorage.com>
From: Ralf Jackson <ralf@globalstorage.com>
Subject: Sort these photos

Hi Pat,

I'm putting together a presentation and I want to use a bunch of photographs of signs that I've been randomly taking over the last couple of years. Could you look through my "Random Photos" album inside the "Ralf Jackson's Album" and move all the photographs of signs to the empty Funny Signs album I made?

Thanks,
Ralf

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Ski trip photos

Hey Pat,

The professional photographer from the company ski trip finally sent me photographs and I want to send them out as part of my weekly email to employees and family members. Please create a new album in the Global Storage albums and put the photographs in it.

I'll send out the email about the photos as soon as you tell me they are up. Don't change any of the titles, in my newsletter I'm going to ask everyone to open the album and create their own titles for the photographs. I don't know if that just works or not.

After you upload the photos can you make sure that none are sideways? Also, make sure there aren't any photos of alcohol or anyone drinking. Susie in marketing may use these later and for some reason she thinks pictures of people drinking are a good thing to send out in a family friendly newsletter, but I don't.

This is going to be great,
Gerald
(The Boss)

The ski trip photographs are on your Desktop in a folder labeled Ski Trip.

To: Pat Jones <pat@globalstorage.com>

From: Gerald Fredricks <gerald@globalstorage.com>

Subject: Put these on your photo site

Hey Pat,

I was just reviewing the ski pictures and I noticed a photo of what looks like alcohol. Please remove it. I don't want any alcohol pictures in this album.

Thanks,

Gerald

(The Boss)

To: Pat Jones <pat@globalstorage.com>

From: Gerald Fredricks <gerald@globalstorage.com>

Subject: Put these on your photo site

Hey Pat,

I was looking through the ski photographs when I noticed one that was sideways. Please go make sure they are all straight. I don't like untidy photo albums.

Thanks,

Gerald

(The Boss)

To: Pat Jones <pat@globalstorage.com>

From: Gerald Fredricks <gerald@globalstorage.com>

Subject: Publicity photos

Hi Pat,

This last week we had a public show case of our new product line. I created an album entitled "New Products" of all the great photographs I collected from the event. But I'm not ready to go public with it yet and really don't want anyone but coworkers seeing it. Could you go through and clean things up a bit? All the photos need to have titles. You can pick whatever title you think is appropriate. I already went through and organized them so everything is in the correct order. I had some trouble because Susan in marketing couldn't see or edit the photographs but I fixed that one.

Thanks,

Gerald

(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Publicity photos

Hi Pat,

I just looked through the New Products album and I found a photograph that was sideways. Please make sure they are all oriented correctly.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Publicity photos

Hi Pat,

I looked through the New Products album and noticed that some of the photographs still have names like IMG123. Could you please give them English sounding titles. The titles don't have to be complex they can be things like "Examining new product."

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: RE: Ski trip photos

Hey Pat,

Some of the Ski Trip photos appear to be sideways. Please fix this.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: RE: Ski trip photos

Hey Pat,

There are still some photographs with titles which clearly mention alcohol. Please change these to some other appropriate title.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: RE: Ski trip photos

Hey Pat,

Since the ski trip photographs will be in the newsletter I want to be sure that they are visible to friends and family. One of the admins claims to have fixed it so that the photos are visible to each employees friends and family. Can you tell me if these photos are visible to your friends and family?

Thanks,
Gerald
(The Boss)

Information: Adventures

Despite having a normal desk job you really like to go out and do fun things on the weekends. When it comes to exciting activities like sky diving you will try anything once. You make sure to post photos of all your adventures so your friends can see. However, your mother is one of those people who panics easily and you know if she ever saw a photograph of you diving out of an airplane you would never hear the end of it. So you make sure not to mention some of your more exciting adventures.

Unlike your work, your friends all put their photos in there own albums.

Say “Done” when you are finished reading this page.

Information

It is now Sunday and you had the weekend off. You are now at your home computer checking email.

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: New photos

Yo Pat,

Here are the better photos from the Building Jumping trip last weekend. Could you put them up on Gallery for me? Just set it up in your album (Pat Jones’s Albums) where everyone already knows to look. Also could you title the photos with the people in them? I had the red parachute, George had the green one and of course your’s was blue.

When you are finished let me know so I can have all our friends go look at it.

Thanks,
Josh

The photos Josh sent are in a folder labeled *Building Jumping* on your desktop.

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: RE: New photos

Hi Pat,

I'm not going to upload these photos because I don't have the time. Please upload them.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: RE: New photos

Hi Pat,

I see the photos are up but they don't have any titles. Please title the photos with the people in them? I had the red parachute, George had the green one and of course your's was blue.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: Are you ok?

Pat,

Are you all right? Are you ok?

I just sent Aunt Sue a link to Jennifer's Baby pictures and she sent me back this photo of you jumping off a building. A BUILDING! Are you crazy? What were you thinking? Do you realize how dangerous what you are doing is? People die from this!

Uncle David already thinks I'm a poor mother, if he sees these photographs I will NEVER hear the end of it. And he is going to be looking as soon as he gets home because I already sent him a link to Jennifer's Baby pictures. What were you thinking? How could you do this to me?

Please, please make sure no more of our family see these photographs.

Mom

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

I'm so glad you made it to the after dark safari trip last month. Wasn't it cool seeing some of those animals at night?

I added some of the photos from the safari trip to my album but somehow most of them turned out sideways. Could you rotate them? Also, I seemed to have uploaded a photo from one of the Pirates games and can't seem to find how to delete it. Could you delete it for me while you are at it?

When you are done let me know so I can email the link out to all our friends. I can't wait for them to see some of these great shots.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

I just checked and some of the photographs from the Safari trip are still sideways.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

That photograph of the Pittsburgh Pirates is still in the Safari album. Could you please help me delete it?

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: Josh Needen <josh@hotmail.com>
Subject: Safari Photos

Yo,

Hey, it just occurred to me. Can our friends even see the Safari album? I'm not sure how to check and I don't want to send it out if they can't see it.

Thanks,
Josh

To: Pat Jones <pat@jones.com>
From: George Wilson <george@hotmail.com>
Subject: Safari Photos

Hey Pat,

I helped Rich re-build his porch this weekend and I took photographs of the whole thing. I think we did an awesome job but I'll let you judge for yourself.

We ripped out half the porch and put in all new mini foundation pieces for each support. Then we put in a whole new frame and put the surface boards back on and painted.

We thought we were done but then Kerry came out, and well you know how much she likes plants, and Rich and I had all these tools laying around. . . So we went ahead and built her a set of planter boxes for the porch. We started out with some small ones then we built some longer ones too.

Finally, we put in some new stairs that will look better and creak less.

Anyway, the whole reason I'm emailing you is that I can't seem to figure out how to put the photos in order. It looks silly right now with the photos of planter boxes appearing before the photos of us putting in the deck. Could you organize them for me?

Thanks,
George

To: Pat Jones <pat@jones.com>
From: Lisa Williams <josh@hotmail.com>
Subject: New photos

Hi Pat,

Thanks for setting up that great album for our Building Jumping trip. I took some photos too so I went ahead and uploaded them. Could you double check that I didn't mess anything up and that all the photos look ok, your photo sharing system always confuses me.

Thanks,
Lisa

To: Pat Jones <pat@jones.com>
From: Lisa Williams <josh@hotmail.com>
Subject: RE: New photos

I just sent out the link to the building jumping photos but I'm getting complaints because our friends can't see the new photographs. Josh sent me an unflattering email about how I shouldn't be allowed to upload photos. What did I do wrong? Could you please fix it?

Thanks,
Lisa

To: Pat Jones <pat@jones.com>
From: Lisa Williams <josh@hotmail.com>
Subject: RE: New photos

I just sent out the link to the building jumping photos but now Josh is making fun of me because one of the photos is sideways. Is there an easy way to turn it back round? Could you please fix it?

Thanks,
Lisa

Information: Pat's Family

Your parents can barely operate their computer much less manage a photo site. So you let your family post photographs in their own album but you help out by checking each album to make sure it is not visible to everyone on the Internet.

You help your parents manage their photos when they upload new albums. Your mother doesn't understand the photo management software on her computer and tends to make a ton of silly mistakes like once accidentally titling your Dad *Fido*. She is perfectionist and not being able to make her photos look perfect really annoys her so you help her out by fixing up the photographs before she lets her friends and family see anything.

Your mother's name is Samantha and all her photographs can be found in "Samantha Jones's Albums".

Say "Done" when you are finished reading this page.

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: New albums

Hi Pat,

I just uploaded the Christmas photos at Jennifer's to my web album. Aunt Sue has been asking about the Christmas photos for months. I'm so glad I finally found time to do this.

I followed the instructions you gave me last time you showed me how to put photos on your photo site but they were so complex I didn't get through all of them. I'm concerned I might have made a few mistakes. To begin with I think I uploaded some photos from my Mexico vacation into the Christmas album. So could you please go and delete any photos that look out of place. Also, I think I might have mixed up a few titles.

Could you please go look at the albums and fix any mistakes I might have made? Let me know when you are done so I can email the family so they can see the pictures.

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: New albums

Hi Pat,

Thanks for fixing up my new albums. I took a quick glance over them and I think there may still be some errors with the titles. The picture with little Henry holding his pillow at Christmas is still labeled *Susan and new pillow*.

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: New albums

Hi Pat,

Thanks for fixing up my new albums. I took a quick glance over them and I think there may still be some problems. The picture of Susan with her arms out is sideways. Could you please make it straight?

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Mom <samantha@jones.com>
Subject: Re: New albums

Hi Pat,

Aunt Sue just emailed me and she says she can't see my Christmas photographs. Where did they go? Why can't she see them?

Thanks,
Mom

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Separate out some photos

Hi Pat,

I uploaded photos of the snow storm, that Pirates game I went too, and a trip to the Phipps Conservatory during their gargoyles exhibit, into that Misc album you created for me. I even managed to create three new albums for the photographs. The only problem is that I can't seem to get the photos moved from the Misc album to the albums they need to be in.

Thanks,
Jennifer

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Baby photos

Hi Pat,

I just took a bunch of photos of my new baby Angela and I want to share the photos with family, friends and coworkers. Could you create a new album for them in "Jennifer Smith's Albums" and put the new photos in it? When you are done I need you to find the cutest one and make it the album cover.

Thanks,
Jennifer

The photos Jennifer sent are in a folder labeled *Angela* on your desktop.

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Re: New albums

Hi Pat,

Mom was all upset about her photos not being quite right so I had her log in for me and tried to fix them myself. But Mom hated all my changes and wants things back the way they were. Could you go back through her Christmas album and just put everything back the way it was?

Thanks,
Jennifer

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Re: New albums

Hi Pat,

Mom says that all the photos used to be straight and now one is not. She isn't letting me touch the computer anymore, can you please fix it.

Thanks,
Jennifer

To: Pat Jones <pat@jones.com>
From: Jennifer Smith <jennifer@hotmail.com>
Subject: Re: New albums

Pat,

Mom is terribly worried that other people not in our family are looking at her photos. I told her that it was fine but could you please just check.

Thanks,
Jennifer

B.2 Online survey

Gallery Prox Info Display (June 27, 2011)

New Page

1. User ID

Page One

2. Did you find working with Gallery today to be: *

- ☐ Very Enjoyable
- ☐ Enjoyable
- ☐ Neutral
- ☐ Unpleasant
- ☐ Very Unpleasant

New Page

Gallery uses a set of icons to indicate information about privacy settings. For each icon below describe what you think the icon means.

3. ☐ *

4. ☐ *

5. ☐ *

6. ☐ *

7. ☐ *

Funny Signs

8. Ralf Jackson asked you to move signs from his "Random Photos" album to another album called "Funny Signs". What was the privacy policy for the Funny Signs album when you left it? *

	True	False	Not Sure
Everybody can add to the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Funny Signs album. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

9. For the answers you marked true or false, how confident are you in your answer? *

Very Confident Confident Neutral Uncertain Very Uncertain Not Applicable

☐ ☐ ☐ ☐ ☐ ☐

Ski Trip

10. Your boss asked you to create an album for the company ski trip. What was the privacy policy for the Ski Trip album when you left it? *

	True	False	Not Sure
Everybody can add to the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Coworkers can view the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Ski Trip album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

11. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

New Products

12. Your boss asked you to review the New Products album for errors. What was the privacy policy for the New Products album when you left it? *

	True	False	Not Sure
Everybody can add to the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the New Products album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

13. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Building Jumping

14. Your friend, Josh, sent you some Building Jumping photos and asked you to create an album. What was the privacy policy for the Building Jumping album when you left it? *

	True	False	Not Sure
Everybody can add to the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Family can add to the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Building Jumping album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

15. For the answers you marked true or false, how confident are you in your answer? *

Very Confident ☐
 Confident ☐
 Neutral ☐
 Uncertain ☐
 Very Uncertain ☐
 Not Applicable ☐

Safari

16. Your friend Josh Needen asked you to rotate some photos in his Safari album. What was the privacy policy for the Safari album when you left it? *

	True	False	Not Sure
Everybody can add to the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Safari album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

17. For the answers you marked true or false, how confident are you in your answer? *

Very Confident ☐
 Confident ☐
 Neutral ☐
 Uncertain ☐
 Very Uncertain ☐
 Not Applicable ☐

Porch

18. Your friend George Wilson asked you to organize his porch building photos; What was the privacy policy for the Porch Building album when you left it? *

	True	False	Not Sure
Everybody can add to the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Porch Building album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

19. For the answers you marked true or false, how confident are you in your answer? *

Very Confident ☐
 Confident ☐
 Neutral ☐
 Uncertain ☐
 Very Uncertain ☐
 Not Applicable ☐

Christmas

20. Your mother asked you to review her Christmas album for errors. What was the privacy policy for the Christmas album when you left it? *

	True	False	Not Sure
Everybody can add to the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Christmas album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

21. For the answers you marked true or false, how confident are you in your answer? *

Very Confident ☐
 Confident ☐
 Neutral ☐
 Uncertain ☐
 Very Uncertain ☐
 Not Applicable ☐

Baby

22. Your sister asked you to create a new album for her baby photos. What was the privacy policy for the Baby Photo album when you left it? *

	True	False	Not Sure
Everybody can add to the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Baby Photo album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

23. For the answers you marked true or false, how confident are you in your answer? *

Very Confident Confident Neutral Uncertain Very Uncertain Not Applicable

☐ ☐ ☐ ☐ ☐ ☐

Misc

24. Your sister asked you to sort some photos from her Misc album to three other albums. What was the privacy policy for the Misc album when you left it? *

	True	False	Not Sure
Everybody can add to the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Misc album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

25. For the answers you marked true or false, how confident are you in your answer? *

Very Confident Confident Neutral Uncertain Very Uncertain Not Applicable

☐ ☐ ☐ ☐ ☐ ☐

Conference

26. In the Global Storage Shared Albums there is an album called Conference. What was the privacy policy for the Conference album when you left it? *

	True	False	Not Sure
Everybody can add to the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Conference album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

27. For the answers you marked true or false, how confident are you in your answer? *

Very Confident Confident Neutral Uncertain Very Uncertain Not Applicable

☐ ☐ ☐ ☐ ☐ ☐

New Desk

28. In your friend George Willson's Albums there is an album called New Desk. What was the privacy policy for the New Desk album when you left it? *

	True	False	Not Sure
Everybody can add to the New Desk album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the New Desk album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the New Desk album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the New Desk album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the New Desk album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the New Desk album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the New Desk album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Friends can **view** the New Desk album.

☐☐☐

29. For the answers you marked true or false, how confident are you in your answer? *

Very
Confident
☐

Confident
☐

Neutral
☐

Uncertain
☐

Very
Uncertain
☐

Not
Applicable
☐

Sky Diving

30. You have an album called Sky Diving. What was the privacy policy for the Sky Diving album when you left it? *

	True	False	Not Sure
Everybody can add to the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can add to the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Sky Diving album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

31. For the answers you marked true or false, how confident are you in your answer? *

Very
Confident
☐

Confident
☐

Neutral
☐

Uncertain
☐

Very
Uncertain
☐

Not
Applicable
☐

Angela's Wedding

32. In your mother's (Samantha Jones) albums there is an album called Angela's Wedding. What was the privacy policy for the Angela's Wedding album when you left it? *

	True	False	Not Sure
Everybody can add to the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Everybody can view the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can add to the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Family can view the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Coworkers can add to the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Coworkers can view the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can add to the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Friends can view the Angela's Wedding album.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

33. For the answers you marked true or false, how confident are you in your answer? *

Very Confident	Confident	Neutral	Uncertain	Very Uncertain	Not Applicable
<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

New Page

34. Which of the following photo sharing applications have you used to share photos?

- ☐ Facebook
- ☐ Flickr
- ☐ Kodak
- ☐ Picasa
- ☐ Photie
- ☐ Photobucket
- ☐ Shutterfly
- ☐ SmugMug
- ☐ Webshots
- ☐ Zoomr

35. How often do you upload photos to online photo sharing sites? *

- ☐ Daily
- ☐ Weekly
- ☐ Monthly
- ☐ Yearly
- ☐ I never upload photographs

36. Which groups do you regularly share photographs with?

- ☐ Co-workers
- ☐ Family
- ☐ Friends
- ☐ Private (visible only to you)
- ☐ Public (visible to anyone on the Internet)

New Page

37. Gender *

- ☐ Male
- ☐ Female

38. What is your age? *

39. What is the highest degree you have received? *

- ☐ 12th grade or less
- ☐ Graduated high school or equivalent
- ☐ Some college, no degree
- ☐ Associate degree
- ☐ Bachelor's degree
- ☐ Post-graduate degree

40. What is your occupation? *

- ☐ Administrative Support (e.g., secretary, assistant)
- ☐ Art, Writing, and Journalism (e.g., author, reporter, sculptor)
- ☐ Business, Management, and Financial (e.g., manager, accountant, banker)

- ☐ Education (e.g., teacher, professor)
- ☐ Legal (e.g., lawyer, law clerk)
- ☐ Medical (e.g., doctor, nurse, dentist)
- ☐ Science, Engineering, IT professional (e.g., researcher, programmer, IT consultant)
- ☐ Service (e.g., retail clerks, server)
- ☐ Skilled Labor (e.g., electrician, plumber, carpenter)
- ☐ Not currently working/Currently unemployed
- ☐ Retired
- ☐ Decline to answer
- ☐ Student (Please specify area of study)
- ☐ Other (Please specify)

Thank You!

Thank you for taking our survey. Your response is very important to us.

Appendix C

Lab study (study 3)

C.1 Printed instructions and emails

In the lab study (study 3) participants were given instructions and emails by the researcher on printed sheets of paper. The remainder of this appendix section is all the instructions and emails used in the study. Each instruction or email was printed on its own sheet of paper, but in the interests of saving space, we show only the textual content of the pages. Each box of text was printed on a single page, without the black border.

The pages which give instructions and the pages with emails that initiate tasks were given to all participants. The pages with emails used to prompt the participant, were given to the participant only if the participant did not complete all parts of the task. The first 14 emails were given to the participant in the order they are depicted here. The remaining emails were presented in a random order.

Instructions

Your name is Pat Jones. You are an administrative assistant at a large company called Global Storage. Global Storage has a company wide photo website, called Gallery, where company related photographs are posted.

Today I will give you emails written on pieces of paper. If you would like to respond to an email just say out loud what you would email back or if you don't want to respond just say "done". Once you have responded I will hand you another piece of paper with the next email.

We are interested in how you approach and solve the issues presented to you. Remember, we are testing the software and how it supports your work with photographs. We are not testing you.

Instructions

Your boss, Gerald, has put you in charge of maintaining the company's online photo website called Gallery. Employees enjoy using this website to share photos amongst themselves and with their family members. Global Storage also uses this website for displaying professional company related photographs.

Many people email you every day asking for you to help them complete photograph management tasks. It is your job to help them but violations of Gerald's photograph policy are not permitted and Gerald has asked you to make any changes necessary to enforce it.

Gerald's photograph policy

1. No photographs containing drugs, alcohol, or anything inappropriate.
2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.
3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.
4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.
5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

To: Pat Jones <pat@globalstorage.com>
From: Charles Taylor <charles@globalstorage.com>
Subject: My new baby

Hi Pat,

Everybody at work has been asking about Dian and my new baby so I took some photographs and posted them to Gallery. Isn't she so cute! Unfortunately, I'm not very good at using Gallery and may have messed a few things up.

The album is called "Charles and Dian's new baby Kerry." The photograph of the card from Dian's mother is sideways and the title has a misspelling that needs to be fixed. Also, I think I accidentally uploaded a photograph of our dog Fido. Could you please delete the photo of Fido?

Thanks,
Charles Taylor

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: People

Hi Pat,

The *People* album has photographs of some of the great people who work here at Global Storage. I think it is wonderful that we have a way to show off some of our employees.

I noticed that Ralf's photograph is in the wrong album. Apparently he put it in "Ralf's Random Photos" album but the prior administrative assistant never moved it to the "People" album. Could you please do so. Also, someone must have thought it would be funny to have a cat as the album cover for the "People" album. Please select some other photograph to be the cover.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Charles' new baby

Hi Pat,

I just saw the photographs of Charles' new baby. I don't know if you noticed but the photograph of the card is sideways. You can tell from the words printed on the card which are sideways.

Please fix it.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Charles' new baby

Hi Pat,

Charles and Dian's new baby is adorable. But the card photograph is titled "Card-ddd" which is not how "Card" is spelled. I'm counting on you to find and fix problems like this in the albums on Gallery.

Please fix it.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Charles' new baby

Hi Pat,

I'm so happy that Charles decided to share the photographs of his new baby with us. However, when I checked the permissions I was disappointed to discover that Everybody on the Internet can see these photographs. I expect you to help employees find and fix problems like this.

Please fix the permissions so they match my policy.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: People

Hi Pat,

Ralph's photograph is still not in the People album. Please move it from the "Ralf's Random Photos" album to the "People" album.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: People

Hi Pat,

The cover of the People album is still a cat. Could you please make the cover be the photograph of Christine.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: What is Gerald's policy?

Hi Pat,

I'm trying to better understand Gerald's policy. Is it ok for me to put the attached photograph from the panel discussion Global Storage hosted on Gallery? If so is there anything I need to make sure to do?

Thanks,
Angela Sebastian



To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: What is Gerald's policy?

Hi Pat,

Sorry to bother you about this again but is it ok for me to put the attached photograph of me trying on wedding dresses on Gallery? If so is there anything I need to make sure to do?

Thanks,
Angela Sebastian



To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: What is Gerald's policy?

Hi Pat,

Last time I bother you about this, I promise. But I also have some photographs from my Bachelorette party. Would the photograph below be ok to post on Gallery?

Thanks,
Angela Sebastian



To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: What is Gerald's policy

Hi Pat,

Does Gerald care what the privacy settings are? Can I just set them up any way I want?

Thanks,
Angela Sebastian

To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: What is Gerald's policy

Hi Pat,

While I was waiting for your email Gerald stopped by and I just asked him what his policy is. I think you may be slightly wrong about what he wants. I've included the policy he told me below.

Thanks,
Angela Sebastian

1. No photographs containing drugs, alcohol, or anything inappropriate.
2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.
3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.
4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.
5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

To: Pat Jones <pat@globalstorage.com>
From: Ralf Jackson <ralf@globalstorage.com>
Subject: Project Competition

Hi Pat,

I was looking at some of our old photographs and I noticed some problems with an album the prior administrative assistant created for me. As you know, Global Storage occasionally hosts college project competitions to help find new talent and to show off how great a company we are.

Could you look through the “Project Competition (2009)” album and fix the errors the last administrative assistant made? All the photographs need to be straight . Also, many of the photographs appear to be duplicates with different titles. Please delete any duplicates.

Thanks,
Ralf

To: Pat Jones <pat@globalstorage.com>
From: Ralf Jackson <ralf@globalstorage.com>
Subject: Funny signs

Hi Pat,

I’m putting together a presentation and I am going to use a bunch of photographs of signs that I’ve been randomly taking over the last couple of years. I know my random photos album isn’t very organized I just like to keep it around so other employees can use some of these random photographs in presentations.

Could you look through the “Ralf’s Random Photos” album and move all the photographs of signs to the empty “Funny Signs” album I made?

Thanks,
Ralf

To: Pat Jones <pat@globalstorage.com>
From: Susie Carol <Susie@globalstorage.com>
Subject: New products presentation

Hi Pat,

This last week we had a public show case of our new product line. I created an album entitled “New Products” of all the great photographs I collected from the event.

Could you go through the “New Products” album I just made and clean things up a bit? All the photos need to have titles. You can pick whatever title you think is appropriate. I already went through and organized them so everything is in the correct order.

Thanks,
Susie Carol,
Global Storage Marketing

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Ski trip photos

Hello Pat,

Remember that great ski trip we took together last month? The ski resort photographer finally sent me photographs and they look great. Your friend Daniel looks hilarious fallen over in the snow, I’m sure it is going to take him a while to live that down.

I created an album called “Pat and Ralf’s ski trip”. Could you please make sure that none of the photos are sideways? Don’t worry about changing any of the titles, I already took care of that. Also, can you pick a more exciting cover photograph?

Thanks,
Josh Needam

To: Pat Jones <pat@globalstorage.com>
From: Gerald Fredricks <gerald@globalstorage.com>
Subject: Re: Ski Trip

Hi Pat,

I just thought I would check up on how you are doing so I checked the photographs Ralf sent out of your ski trip. I think I need to remind you of my policy about “acceptable” photograph albums.

1. No photographs containing drugs, alcohol, or anything inappropriate.
2. Personal photos from trips or events not related to work are ok but should only be visible to employees and their families.
3. Professional photographs that involve Global Storage need to be visible to everybody on the Internet so everybody can see how great of a company we are.
4. It is ok for Global Storage employees to add or edit photographs but it isn't ok for anyone else.
5. No photographs that are sideways, have misspellings, duplicated, or excessively blurry.

Thanks,
Gerald
(The Boss)

To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: Conference venue photographs

Hi Pat,

As you know, Global Storage just finished hosting a small conference called BoxTalk and I'm trying to get all the venue photographs posted on Gallery so I can put a link to them on the public website.

You can find the photographs in the "BoxTalk Venue" album. I may have uploaded some of the photographs multiple times so if you see any duplicates feel free to delete them. Also I don't like the current album cover, please select a different photograph and make it the cover. If you see anything else wrong go ahead and fix it.

Thanks,
Angela Sebastian

To: Pat Jones <pat@globalstorage.com>
From: Angela Sebastian <angela@globalstorage.com>
Subject: Conference panel discussion photographs

Hi Pat,

We had a great panel discussion at the BoxTalk conference Global Storage recently hosted. I had Josh take some photographs of the panel discussion which he put on Gallery for me and I would like to put a link to them on the conference forum.

You can find the photographs in the "BoxTalk Panel Discussion" album. Some of the photographs are in the wrong order. The photos of the sandwiches and Jason standing at the podium all need to be at the beginning. The photographs of the panel attendees standing up need to be at the end.

Thanks,
Angela Sebastian

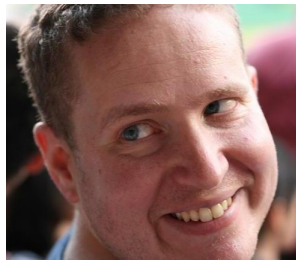
To: Pat Jones <pat@globalstorage.com>
From: Charles Taylor <Charles@globalstorage.com>
Subject: Pirates Game

Hi Pat,

I just realized I still have a great collection of photographs from the trip we took to the Pittsburgh Pirates baseball game that I haven't yet put on the Gallery site. So I thought I had better put them up on the site. Better late than never, right.

I don't actually know who some of these people are so I only titled the people I know. I've put the photographs in an album called "Pittsburgh Pirates". Could you please go and title all the people you recognize?

Thanks,
Charles



William Barish



Chris Macolm



Cathy Keen

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Cool sculptures in Oregon

Hi Pat,

I recently went on a trip to visit my friends in Oregon. They took me to this great event where contestants make moving sculptures and then race them. They have to both race down the road and successfully peddle them up a sand dune. Some of the sculptures are very inventive.

I uploaded them into an album called "Cool Moving Sculptures" and titled some of them. Could you come up with good titles for the rest? Also could you pick your favorite as the cover?

Thanks,
Josh Needam

To: Pat Jones <pat@globalstorage.com>
From: Susie Carol <susie@globalstorage.com>
Subject: Seattle candlelight parade

Hi Pat,

I don't know if you are aware but every year Seattle has a candlelight parade. All the floats have lights on them and the parade happens after dark. This year Global Storage decided to sponsor a float and I took lots of photographs.

Please help me clean up the "Seattle Candlelight Parade (2011)" album. I took lots of great photographs but I'd rather if this album was all on one page. So please delete your least favorite photographs so that there are no more than 12 photos in this album.

Thanks,
Susie Carol
Global Storage Marketing

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Community service

Hi Pat,

The Global Storage Gives Back official community service day this weekend was a big success. Over half of Global Storage's employees decided to participate by doing everything from helping build houses to cleaning up streets. I volunteered with Habitat For Humanity building a porch on a new house. Susie in Marketing asked me to create an album with all the photographs I took at the event so she can use it to show off how great this company is.

I put all of my photographs in a new album called "Global Storage Gives Back". However, some are sideways, please help me out by turning them around straight. Also, I think I uploaded a bunch of other random photographs into the album by accident. Could you please delete any photographs that don't involve building porches.

Thanks,
Josh

To: Pat Jones <pat@globalstorage.com>
From: Charles Taylor <charles@globalstorage.com>
Subject: Grace's Birthday

Hi Pat,

My daughter Grace just had a birthday and I made sure to photograph the whole event and put the photos in a new album. The cake in particular was very nice looking and I got several shots of that. Ya, I know I made it but that doesn't make it any less awesome.

Please look through the "Grace's Birthday" album and just make sure everything looks ok. I may have gone a bit overboard with photographing the cake, go ahead and pick your favorite(s) and delete the rest.

Thanks,
Charles

To: Pat Jones <pat@globalstorage.com>
From: Josh Needam <josh@globalstorage.com>
Subject: Florence Photographs

Hi Pat,

I just got back from my vacation to Florence, Italy. I took photos while exploring the city and now everyone keeps asking me about all the great sights I saw. So I thought I would put together a photo album of Florence.

I put all the photographs in the album “Josh’s trip to Florence Italy.” Could you please help me out by making up titles for the couple of photographs I couldn’t think of good titles for. Also, can you pick your favorite photograph as the album cover? The one I have now is just too generic.

Thanks,
Josh

To: Pat Jones <pat@globalstorage.com>
From: Susie Carol <susie@globalstorage.com>
Subject: Factory Tour

Hi Pat,

In an effort to promote public awareness Global Storage is now offering factory tours at some of our factories. I’ve taken several photographs from the tour at one of our closer factories and put them on Gallery.

Could you go through the “Factory Tour” album and clean up the titles? All the photographs have titles but some of them have dashes in the middle of the title and I don’t want any to have dashes. Also, could you make the photo of people waiting in line be the album cover?

Thanks,
Susie Carol

To: Pat Jones <pat@globalstorage.com>
From: Emma Johnson <emma@globalstorage.com>
Subject: Pumpkin Carving

Hi Pat,

I got a bunch of my girl friends together for some relaxing pumpkin carving fun. We got some great pumpkins that I wanted to show off. I particularly like the one with the witch in the apple (my creation).

Could you please help me fix the order of the photographs in the “Halloween Pumpkin Carving” album? Right now there are photographs of carved pumpkins before the photographs of them being carved. Also, could you make the photo of the pumpkins with the lights out be the album cover?

Thanks,
Emma

Appendix D

Online study (study 4)

D.1 Online survey

Gallery MTurk Questions

Opinion

1. Please indicate if you agree or disagree with the following statements about the **work** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
It was easy to determine if there was an error in the permissions . *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
It was easy to determine if there was a spelling error in the title . *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
It was easy to determine if a photo was sideways . *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
It was easy to determine if there was an error in the tags . *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

2. Please indicate if you agree or disagree with the following statements about the **home** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I identified and corrected all the spelling errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I identified and corrected all the rotation errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I identified and corrected all the tag errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I identified and corrected all the permission errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

3. Please indicate if you agree or disagree with the following statements about the **home** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
It was easy to determine if a photo was sideways . *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
It was easy to determine if there was an error in the permissions . *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
It was easy to determine if there was an error in the tags . *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

It was easy to determine if there was a **spelling** error in the **title**. *

☐
☐
☐
☐
☐

4. Please indicate if you agree or disagree with the following statements about the **work** website. *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I identified and corrected all the rotation errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I identified and corrected all the spelling errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I identified and corrected all the permission errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I identified and corrected all the tag errors. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

Memory

5. For the **White Water Kayaking album** which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view White Water Kayaking Album	_ can currently view White Water Kayaking Album
Everybody on the internet	<input type="checkbox"/>	<input type="checkbox"/>
Adventure Friends	<input type="checkbox"/>	<input type="checkbox"/>
Animal Shelter	<input type="checkbox"/>	<input type="checkbox"/>
Family	<input type="checkbox"/>	<input type="checkbox"/>
Pat Jones	<input type="checkbox"/>	<input type="checkbox"/>
I don't know	<input type="checkbox"/>	<input type="checkbox"/>

6. For the **Teapots album** which of the following groups would Pat's boss want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Teapots Albums	_ can currently view Teapots Album
Everybody on the internet	<input type="checkbox"/>	<input type="checkbox"/>
Dezig Designers	<input type="checkbox"/>	<input type="checkbox"/>

Innovative Teapots	<input type="checkbox"/>	<input type="checkbox"/>
Purse Central	<input type="checkbox"/>	<input type="checkbox"/>
Starlight Phones	<input type="checkbox"/>	<input type="checkbox"/>
I don't know	<input type="checkbox"/>	<input type="checkbox"/>

7. For the **Bags with Toy album** which of the following groups would Pat's boss want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Bag With Toy Album	_ can currently view Bag With Toy Album
Everybody on the internet	<input type="checkbox"/>	<input type="checkbox"/>
Dezig Designers	<input type="checkbox"/>	<input type="checkbox"/>
Innovative Teapots	<input type="checkbox"/>	<input type="checkbox"/>
Purse Central	<input type="checkbox"/>	<input type="checkbox"/>
Starlight Phones	<input type="checkbox"/>	<input type="checkbox"/>
I don't know	<input type="checkbox"/>	<input type="checkbox"/>

8. For the **Animal Shelter album** which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Animal Shelter Album	_ can currently view Animal Shelter Album
Everybody on the internet	<input type="checkbox"/>	<input type="checkbox"/>
Adventure Friends	<input type="checkbox"/>	<input type="checkbox"/>
Animal Shelter	<input type="checkbox"/>	<input type="checkbox"/>
Family	<input type="checkbox"/>	<input type="checkbox"/>
Pat Jones	<input type="checkbox"/>	<input type="checkbox"/>
I don't know	<input type="checkbox"/>	<input type="checkbox"/>

9. For the **Inspirational Phones album** which of the following groups would Pat's boss want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Inspirational Phones Album	_ can currently view Inspirational Phones Album
Everybody on the internet	<input type="checkbox"/>	<input type="checkbox"/>

Dezig Designers	<input type="checkbox"/>	<input type="checkbox"/>
Innovative Teapots	<input type="checkbox"/>	<input type="checkbox"/>
Purse Central	<input type="checkbox"/>	<input type="checkbox"/>
Starlight Phones	<input type="checkbox"/>	<input type="checkbox"/>
I don't know	<input type="checkbox"/>	<input type="checkbox"/>

10. For the **Family Calendar album** which of the following groups would Pat want to be able to view the photos? Which of the following groups can currently view the album based on the current permission settings? *

	Pat wants _ to be able to view Family Calendar Album	_ can currently view Family CalendarAlbum
Everybody on the internet	<input type="checkbox"/>	<input type="checkbox"/>
Adventure Friends	<input type="checkbox"/>	<input type="checkbox"/>
Animal Shelter	<input type="checkbox"/>	<input type="checkbox"/>
Family	<input type="checkbox"/>	<input type="checkbox"/>
Pat Jones	<input type="checkbox"/>	<input type="checkbox"/>
I don't know	<input type="checkbox"/>	<input type="checkbox"/>

Experiences

11. How frequently do you upload and share photographs? *

- ☐ A few times a day
- ☐ A few times a week
- ☐ A few times a month
- ☐ A few times a year
- ☐ Less than once a year
- ☐ Never

12. Which of the following photo sharing sites have you ever used to share photographs? *

- ☐ Flickr
- ☐ Snapfish
- ☐ Photobucket

- ☐ Shutterfly
- ☐ Picasa Web Albums
- ☐ Kodak
- ☐ Phanfare
- ☐ SmugMug
- ☐ Facebook
- ☐ Other

13. Please indicate if you agree or disagree with the following statements *

	Strongly disagree	Disagree	Neutral	Agree	Strongly agree
I am a detail oriented person. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I feel confident about my ability to manage the privacy settings on the photo sharing sites I use. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I feel confident about my ability to manage tags on the photo sharing sites I use. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Most businesses handle the personal information they collect about consumers in a proper and confidential way. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I generally notice whether or not a website I am visiting has a privacy policy. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I am concerned about threats to my personal privacy online today. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
I do not care who sees the photos I post online. *	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

14. Have you ever had a negative experience after sharing a photograph on a photograph sharing site or a social networking site such as Facebook? *

- ☐ Yes
- ☐ No

15. Have you ever done any of the following

- ☐ Created a new group and only shared photos with that group.
- ☐ Changed the privacy settings for a specific photo or album.
- ☐ Set privacy settings on a photo sharing site to "Friends Only."
- ☐ Emailed a photo instead of putting it on a sharing site because of privacy concerns.
- ☐ Other

Demographics

16. Gender *

- ☐ Male
- ☐ Female

17. Age *

18. Select the category that best describes your profession. *

- ☐ Accounting / Finance / Banking
- ☐ Administration / Clerical / Reception
- ☐ Advertisement / PR
- ☐ Architecture / Design
- ☐ Arts/Leisure / Entertainment
- ☐ Beauty / Fashion
- ☐ Buying / Purchasing
- ☐ Construction
- ☐ Consulting
- ☐ Customer Service
- ☐ Distribution

- ☐ Education
 - ☐ Health Care (Physical & Mental)
 - ☐ Human resources management
 - ☐ Management (Senior / Corporate)
 - ☐ News / Information
 - ☐ Operations / Logistics
 - ☐ Planning (Meeting, Events, etc.)
 - ☐ Production
 - ☐ Real Estate
 - ☐ Research
 - ☐ Restaurant / Food service
 - ☐ Sales / Marketing
 - ☐ Science / Technology / Programming
 - ☐ Social service
 - ☐ Student
 - ☐ Other
 - ☐ N/A - Unemployed / Retired / Homemaker
-

19. Highest degree obtained *

- ☐ 12th grade or less
 - ☐ Graduated high school or equivalent
 - ☐ Some college, no degree
 - ☐ Associate degree
 - ☐ Bachelor's degree
 - ☐ Post-graduate degree
-

20. Household income *

- ☐ Less than \$25,000

- ☐ \$25,000 to \$34,999
- ☐ \$35,000 to \$49,999
- ☐ \$50,000 to \$74,999
- ☐ \$75,000 to \$99,999
- ☐ \$100,000 to \$124,999
- ☐ \$125,000 to \$149,999
- ☐ \$150,000 or more

Thank You!

Thank you for completing this study.

Appendix E

Evaluation data

Table E.1: Data depicted in Figure 7.1(a).

Percent of time on page	Number of fixations
0	6
5	5
10	2
15	10
20	5
25	12
30	10
35	16
40	16
45	7
50	6
55	13
60	4
65	6
70	6
75	3
80	9
85	10
90	8
95	19

Number of fixations by under photo participants on proximity display for each 10% of time on page.

Table E.2: Data depicted in Figure 7.1(b).

Percent of time on page	Number of fixations
0	8
5	2
10	4
15	0
20	0
25	3
30	1
35	0
40	2
45	0
50	1
55	3
60	1
65	2
70	7
75	4
80	5
85	6
90	10
95	10

Number of fixations by sidebar participants on proximity display for each 10% of time on page.

Table E.3: Data depicted in Figure 7.2(a).

Difference between number of permissions checked on tasks with errors and without errors	Number of control participants	Number of under photo participants
-1	1	0
0	11	5
1	1	3
2	0	5
3	1	3

Number of tasks where the lab study participants checked permissions and there was an error subtracted by the number of tasks where participants checked permissions and there was not an error (first column). For example, we can see that 11 control condition participants checked the same number of permissions in tasks with errors as they did in tasks without errors (second row).

Table E.4: Data depicted in Figure 7.2(b).

Difference between number of permissions checked on tasks with errors and without errors	Number of control participants	Number of under photo participants
-3	2	0
-2	6	1
-1	15	10
0	82	70
1	21	28
2	4	10
3	1	5
4	0	7

Number of tasks where the online study participants checked permissions and there was an error subtracted by the number of tasks where participants checked permissions and there was not an error (first column).

Table E.5: Data depicted in Figure 7.3(a).

Number of tasks	Number of control participants	Number of under photo participants
0	4	2
1	1	2
2	2	3
3	0	1
4	1	0
5	2	1
6	0	1
7	0	2
8	0	2
9	0	0
10	1	1
11	0	0
12	1	1
13	0	0
14	1	0

Number of tasks where the permission-modification interface was opened by lab participants in the under photo condition. Graphs of other conditions are nearly identical.

Table E.6: Data depicted in Figure 7.3(b).

Number of tasks	Number of control participants	Number of under photo participants
0	62	50
1	17	24
2	9	12
3	5	7
4	6	11
5	8	7
6	10	9
7	6	6
8	8	5

Number of tasks where the permission-modification interface was opened by online participants in the under photo condition. Graphs of other conditions are nearly identical.

Table E.7: Data depicted in Figure 7.4 a.

Number of tasks	Order event engaged in	Event
70	first	cover
73	middle	cover
79	last	cover
11	only	cover
118	first	delete
87	middle	delete
57	last	delete
36	only	delete
39	first	move
58	middle	move
54	last	move
20	first	permissions
32	middle	permissions
101	last	permissions
4	only	permissions
74	first	rename
98	middle	rename
76	last	rename
23	only	rename
94	first	rotate
111	middle	rotate
22	last	rotate
1	only	rotate

While working on tasks in the lab study participants were free to engage in actions in any order, including interleaving actions. For example: participants could rotate a photo, delete a photo, then rotate a photo. This table shows the first time an action of that type was engaged in during a particular task and whether that action was the first action, the last, neither first nor last (middle), or the only action engaged in.

Table E.8: Data depicted in Figure 7.4 b.

Number of tasks	Order event engaged in	Event
79	first	cover
72	middle	cover
71	last	cover
11	only	cover
75	first	delete
122	middle	delete
65	last	delete
36	only	delete
25	first	move
52	middle	move
74	last	move
17	first	permissions
28	middle	permissions
108	last	permissions
4	only	permissions
49	first	rename
120	middle	rename
79	last	rename
23	only	rename
150	first	rotate
61	middle	rotate
16	last	rotate
1	only	rotate

While working on tasks in the lab study participants were free to engage in actions in any order, including interleaving actions. For example: participants could rotate a photo, delete a photo, then rotate a photo. This table shows the last time an action of that type was engaged in during a particular task and whether that action was the first action, the last, neither first nor last (middle), or the only action engaged in.

Table E.9: Data depicted in Figure 7.5 (sidebar).

Number of times the permission modification dialog opened	Seconds into the task (10sec intervals)	Treatment	Condition
1	0	sidebar	control
8	10	sidebar	control
16	20	sidebar	control
13	30	sidebar	control
20	40	sidebar	control
13	50	sidebar	control
19	60	sidebar	control
20	70	sidebar	control
31	80	sidebar	control
29	90	sidebar	control
26	100	sidebar	control
26	110	sidebar	control
15	120	sidebar	control
5	0	sidebar	experimental
15	10	sidebar	experimental
17	20	sidebar	experimental
15	30	sidebar	experimental
13	40	sidebar	experimental
17	50	sidebar	experimental
14	60	sidebar	experimental
20	70	sidebar	experimental
23	80	sidebar	experimental
18	90	sidebar	experimental
26	100	sidebar	experimental
22	110	sidebar	experimental
15	120	sidebar	experimental
1	130	sidebar	experimental

Table E.10: Data depicted in Figure 7.5 (under).

Number of times the permission modification dialog opened	Seconds into the task (10sec intervals)	Treatment	Condition
2	0	under	control
13	10	under	control
27	20	under	control
18	30	under	control
22	40	under	control
17	50	under	control
25	60	under	control
18	70	under	control
38	80	under	control
38	90	under	control
23	100	under	control
22	110	under	control
9	120	under	control
1	130	under	control
4	0	under	experimental
21	10	under	experimental
32	20	under	experimental
23	30	under	experimental
15	40	under	experimental
21	50	under	experimental
20	60	under	experimental
23	70	under	experimental
20	80	under	experimental
33	90	under	experimental
19	100	under	experimental
17	110	under	experimental
17	120	under	experimental
1	130	under	experimental

Table E.11: Data depicted in Figure 7.5 (mixed).

Number of times the permission modification dialog opened	Seconds into the task (10sec intervals)	Treatment	Condition
5	0	mixed	control
10	10	mixed	control
14	20	mixed	control
11	30	mixed	control
16	40	mixed	control
12	50	mixed	control
21	60	mixed	control
20	70	mixed	control
26	80	mixed	control
16	90	mixed	control
23	100	mixed	control
29	110	mixed	control
12	120	mixed	control
3	0	mixed	experimental
7	10	mixed	experimental
27	20	mixed	experimental
25	30	mixed	experimental
20	40	mixed	experimental
13	50	mixed	experimental
14	60	mixed	experimental
16	70	mixed	experimental
16	80	mixed	experimental
17	90	mixed	experimental
17	100	mixed	experimental
15	110	mixed	experimental
6	120	mixed	experimental

Table E.12: Data depicted in Figure 7.5 (facebook).

Number of times the permission modification dialog opened	Seconds into the task (10sec intervals)	Treatment	Condition
3	0	facebook	control
11	10	facebook	control
24	20	facebook	control
24	30	facebook	control
20	40	facebook	control
19	50	facebook	control
28	60	facebook	control
33	70	facebook	control
29	80	facebook	control
29	90	facebook	control
18	100	facebook	control
27	110	facebook	control
18	120	facebook	control
1	130	facebook	control
5	0	facebook	experimental
18	10	facebook	experimental
23	20	facebook	experimental
28	30	facebook	experimental
17	40	facebook	experimental
14	50	facebook	experimental
30	60	facebook	experimental
28	70	facebook	experimental
37	80	facebook	experimental
24	90	facebook	experimental
26	100	facebook	experimental
26	110	facebook	experimental
8	120	facebook	experimental

Table E.13: Data depicted in Figure 7.5 (audit).

Number of times the permission modification dialog opened	Seconds into the task (10sec intervals)	Treatment	Condition
2	0	audit	control
21	10	audit	control
33	20	audit	control
22	30	audit	control
21	40	audit	control
18	50	audit	control
30	60	audit	control
33	70	audit	control
33	80	audit	control
33	90	audit	control
43	100	audit	control
20	110	audit	control
14	120	audit	control
1	0	audit	experimental
10	10	audit	experimental
32	20	audit	experimental
26	30	audit	experimental
20	40	audit	experimental
16	50	audit	experimental
18	60	audit	experimental
27	70	audit	experimental
44	80	audit	experimental
42	90	audit	experimental
33	100	audit	experimental
21	110	audit	experimental
14	120	audit	experimental
1	130	audit	experimental

Table E.14: Data depicted in Figure 7.6.

Order	Alcohol	Personal	Professional	Edit	Sideways
0	0	4	5	27	3
1	35	0	2	0	5
2	5	8	4	2	23
3	2	18	15	5	1
4	0	10	15	5	9
5	0	2	1	3	1

As part of the lab study verbal post-survey, participants were asked to recall Gerald's rules, in their own words. The above table shows the the order (column 1) in which participants recalled each of the rules (columns). Each cell shows the number of participants who recalled that rule in that order position.

Table E.15: Data depicted in Figure 7.7.

Number of action events	Seconds into task (10 second intervals)	Event
31	0	Permission modification
134	10	Permission modification
245	20	Permission modification
205	30	Permission modification
184	40	Permission modification
160	50	Permission modification
219	60	Permission modification
238	70	Permission modification
297	80	Permission modification
279	90	Permission modification
254	100	Permission modification
225	110	Permission modification
128	120	Permission modification
5	130	Permission modification
30	0	Delete
169	10	Delete
606	20	Delete
1406	30	Delete
1778	40	Delete
1670	50	Delete
1449	60	Delete
1270	70	Delete
1009	80	Delete
845	90	Delete
608	100	Delete
427	110	Delete
268	120	Delete
15	130	Delete

Number of seconds into a task that an action was engaged in (10 second intervals). Table shows all participants across all conditions, both with and without permission proximity displays. Events from task 1 and the training are excluded to remove bias caused by prompting participants.

Table E.16: Data depicted in Figure 7.7 cont.

Number of action events	Seconds into task (10 second intervals)	Event
4	0	Organize opened
69	10	Organize opened
285	20	Organize opened
379	30	Organize opened
374	40	Organize opened
413	50	Organize opened
373	60	Organize opened
295	70	Organize opened
241	80	Organize opened
185	90	Organize opened
159	100	Organize opened
103	110	Organize opened
49	120	Organize opened
5	130	Organize opened
148	0	Rename
368	10	Rename
515	20	Rename
1294	30	Rename
1938	40	Rename
2268	50	Rename
2427	60	Rename
2332	70	Rename
2120	80	Rename
1861	90	Rename
1534	100	Rename
1192	110	Rename
695	120	Rename
23	130	Rename

Number of seconds into a task that an action was engaged in (10 second intervals). Table shows all participants across all conditions, both with and without permission proximity displays. Events from task 1 and the training are excluded to remove bias caused by prompting participants.

Table E.17: Data depicted in Figure 7.7 cont.

Number of action events	Seconds into task (10 second intervals)	Event
40	0	Rotate
177	10	Rotate
990	20	Rotate
1858	30	Rotate
2099	40	Rotate
1868	50	Rotate
1471	60	Rotate
1309	70	Rotate
986	80	Rotate
824	90	Rotate
651	100	Rotate
442	110	Rotate
246	120	Rotate
19	130	Rotate
18	0	Tag modification
74	10	Tag modification
179	20	Tag modification
185	30	Tag modification
126	40	Tag modification
139	50	Tag modification
160	60	Tag modification
218	70	Tag modification
239	80	Tag modification
222	90	Tag modification
216	100	Tag modification
187	110	Tag modification
105	120	Tag modification
4	130	Tag modification

Number of seconds into a task that an action was engaged in (10 second intervals). Table shows all participants across all conditions, both with and without permission proximity displays. Events from task 1 and the training are excluded to remove bias caused by prompting participants.

